

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

Mangala N., Murtaza Rangwala, Aishwarya S., B. Eswara Reddy, Rajkumar Buyya,
K. R. Venugopal, S. S. Iyengar, L. M. Patnaik



Abstract—Healthcare has become exceptionally sophisticated, as wearables and connected medical devices are revolutionising 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 extremely important in data-driven healthcare. In this paper, we propose a multi-layer IoT, Edge and Cloud architecture to enhance the speed of response for emergency healthcare by distributing tasks based on response criticality and permanence of storage. Privacy of patient data is assured by proposing a Differential Privacy framework across several machine learning models such as 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 86% 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 $\epsilon = 5.0$, supervised algorithms achieve 82-84% accuracy while reducing attribute inference attacks by up to 18% and data reconstruction correlation by 70%. Blockchain security further 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.

- Mangala N. 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
- Aishwarya S. 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
- K. R. 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
- S. S. Iyengar is a Distinguished Professor and Director of Computer Science, Florida International University, Miami, USA. E-mail: iyengar@cs.fiu.edu
- L. M. Patnaik is Adjunct Professor and NASI Senior Scientist in the National Institute of Advanced Studies, Bangalore, India. E-mail: lalit-blr@gmail.com

1 INTRODUCTION

The Internet of Things (IoT) has become an integral part of the modern application ecosystem, transforming not only organisational and industrial operations but also our daily lives. In the healthcare sector, digitalisation 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. These Healthcare IoT devices require configuration with the user's personal information such as age, gender and location. Healthcare Analytics is a branch of data analytics that harnesses this patient data to generate descriptive, diagnostic, predictive, and prescriptive insights, fundamentally supporting evidence-based medical decision-making. Moreover, through the application of AI/ML techniques, 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, with adversaries potentially exploiting patient information for blackmail, ransom demands, manipulation, generating artificial market demand, and targeted marketing campaigns [2]. The protection of patients' data thus becomes 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 organisations 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, specifically, establishes national standards governing the security of health information during transfer, reception, processing and sharing [5]. Despite these regulatory safeguards, the persistence of data breaches underscores 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 [6].

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 privacy of the patient because combinations of other fields such as postcode and date of birth can combine uniquely to re-identify the person.

A significant case illustrating privacy vulnerabilities in healthcare data occurred when Sweeney [7] demonstrated how medical records could be re-identified despite anonymisation. In one instance, she successfully re-identified the medical records of the Governor of Massachusetts by executing merely three queries on an anonymised hospital dataset released by the Massachusetts Group Insurance Commission (GIC). The re-identification was facilitated by cross-referencing public information—news reports of the governor’s collapse on 18 May 1996 and subsequent hospitalisation—with voter registration records containing demographic identifiers (names, addresses, postcodes, birth dates, and gender). 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 [8], 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 to safeguard against such privacy issues.

The *Differential Privacy guarantee* states that in a database, adding or dropping a row does not affect the output of the analysis, and an adversary cannot tell if 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 [9] and multi-party secure computing [10] enable operations on protected

data while maintaining confidentiality, these methods frequently present substantial computational complexity challenges [11]. 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

The primary focus of this research is to enhance 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 [12], [13] and blockchain-based healthcare systems [3], [14]–[16] 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, that is, 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 mech-

anism 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 followed by conclusions and future directions in Section 9.

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 optimisation, genomic analysis, and surgical automation [8]. 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 [17], Random Forests [18], and Naive Bayes [19]. 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 [20] 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 generalisation, 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 categorisation. 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 [21]. The DP guarantee for two datasets is given by:

$$\Pr[M(D) \in S] \leq e^\epsilon \Pr[M(D') \in S]$$

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 [22] 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 [23].

(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 L_2 -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 epsilon parameter quantifies the strength of privacy protection. A smaller epsilon 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 [24].

(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 randomisation 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 [25]. 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:

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

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 scale of the Gaussian noise satisfies:

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

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 L_2 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 L_2 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:

- Healthcare*: DP enables secure sharing of electronic health records and clinical trial data, allowing researchers to extract insights whilst maintaining patient confidentiality [26], [27].
- 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 [28].

(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 [29], [30].

(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 [31].

(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 [32].

(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 [33].

(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 [34].

(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 [35].

(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 [30].

(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 [36].

2.2.7 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 [27].

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 hospitalisations and treatments, can be traced and categorised 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 [37]. 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.

2.3 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.

- 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 [38]. In healthcare contexts, 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.
- Cloud Computing*: Cloud computing provides on-demand network access to a shared pool of configurable computing resources, including servers, storage,

TABLE 1: Applications of Differential Privacy

Application	Description	Factors to Consider
Data Analysis	Aggregating and analyzing sensitive data	<ul style="list-style-type: none"> - 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	<ul style="list-style-type: none"> - Protecting user preferences while providing personalized suggestions - Trade-off between utility and privacy - Ensuring diversity in recommendations
Healthcare Research	Medical data analysis and research	<ul style="list-style-type: none"> - 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	<ul style="list-style-type: none"> - 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	<ul style="list-style-type: none"> - 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	<ul style="list-style-type: none"> - 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	<ul style="list-style-type: none"> - 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	<ul style="list-style-type: none"> - 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	<ul style="list-style-type: none"> - 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	<ul style="list-style-type: none"> - Protecting user privacy while analyzing traffic data - Noise addition to GPS data for anonymity - Data aggregation techniques for traffic analysis - Privacy-aware routing algorithms

applications and databases [39]. 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 organisations 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. [40] 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. [12] 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. [13] 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. [41] 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. [42] 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. [43] described a DP framework utilising a noisy max algorithm with Laplace mechanism that provided accuracy guarantees, suggesting potential for practical application with extensive empirical evaluation on real-world datasets.

Foundational work by Dwork [44] 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. [45] 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.

TABLE 2: DP Literature Review

Author Concept/Model	Algorithm Implementation	Performance Advantages	Research Gaps Future Work
Huang et al., [46] 2020, DP-ADMM:ADMM-Based Distributed Learning with DP	- DP-Alternating direction method of multipliers (DP-ADMM) - Real-world data set: 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., [33] 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., [13] 2020, DP-Based Blockchain for Industrial Internet-of-Things	- Blockchain Internet of Edge model - Q-learning algorithm - Laplace's distribution noise	- BloE model enhance the Privacy-Preserving capacity - Tamper resistant	Study impact of data and noise on energy cost in future works
Li et al., [47] 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., [25] 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 tradeoff between data utility and privacy leaks - Reduce Data correlation	Data correlation may bring new errors with different queries
Bi Jia et al., [12], 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., [48] 2019, 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., [49] 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., [50] 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., [51], 2020, 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

4 PROBLEM STATEMENT

Medical records containing patients' 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 this data is securely transmitted to cloud storage for long-term preservation, diagnosis and analytics.

Objectives:

- (i) Enhance the response speed of IoT-Cloud healthcare systems
- (ii) Provide robust privacy guarantees for patients' 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 optimisation within the proposed frame-

work 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 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 and Cloud computing architecture to enhance communication efficiency, response time, computational capacity, security mechanisms and storage capabilities
- (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.

Privacy protection is achieved through DP techniques applied at 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 e^ϵ , ensuring that, healthcare institutions can share data with researchers whilst ensuring individual patients cannot be re-identified, 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.




Fig. 3: Multi-Layered Compute-Storage Architecture

This layered architecture directly addresses the availability and resilience requirements identified in Section 5. By dis-

tributing processing across multiple tiers, the system maintains emergency response capabilities even if higher-layer Edge or Cloud nodes are compromised or unavailable. The 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 epsilon (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. Our permissioned blockchain with Byzantine fault tolerance ensures integrity even with up to $f < n/3$ compromised consensus nodes, addressing the consensus attack vectors identified in the threat model.

(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).

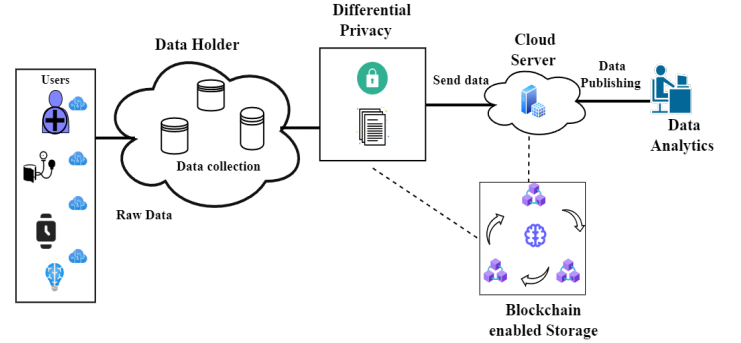


Fig. 4: Securing Healthcare Data by DP and Blockchain

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 analysed 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.

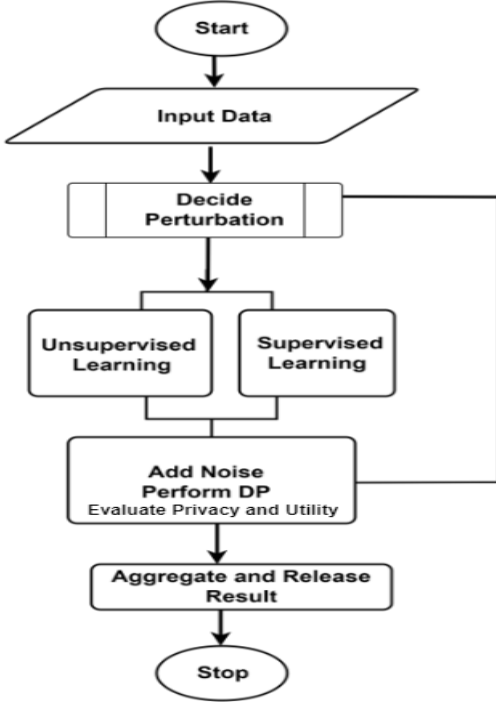


Fig. 5: DP Mechanism with ML

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 $\sigma = \frac{\Delta f}{\epsilon}$. 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.

Algorithm 1: Differentially Private Random Forest

Data: Training data X , labels y , privacy budget ϵ , sensitivity Δ

Result: Differentially private random forest DP_{RF}

```

1 foreach tree  $t \in RandomForest$  do
2   Train base decision tree:  $t \leftarrow \text{TRAINTREE}(X, y)$ ;
3   foreach leaf  $l \in t$  do
4      $\sigma \leftarrow \Delta/\epsilon$ ;
5      $noise \leftarrow \text{SampleGaussian}(0, \sigma)$ ;
6      $l.\text{prediction} \leftarrow l.\text{prediction} + noise$ ;
7   end
8 end
9 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 $\text{Lap}(\Delta f/\epsilon)$ 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

```

1  $C \leftarrow \text{INITIALISECENTROIDS}(D, k)$ ;
2 for  $t = 1$  to  $T$  do
3    $S \leftarrow \text{AssignClusters}(D, C)$ ;
4   foreach cluster  $c_i \in C$  do
5      $\tilde{\mu}_i \leftarrow \text{mean}(S_i) + \text{Lap}(\Delta/\epsilon)$ ;
6      $C \leftarrow C \cup \tilde{\mu}_i$ ;
7   end
8 end
9 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 optimisation 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 $\eta \sim \text{Lap}(\Delta/\epsilon)$ 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.

Algorithm 3: Differentially Private Logistic Regression

Data: Training set D , ϵ , Δ , learning rate η , iterations T

Result: Private weights w

```

1 Initialise  $w \leftarrow \mathbf{0}$ ;
2 for  $t = 1$  to  $T$  do
3   Compute  $\nabla L(w) = \sum_{(x_i, y_i) \in D} (p_i - y_i)x_i$ ;
4    $\nabla L(w) \leftarrow \text{PerturbGradient}(\nabla L(w), \Delta, \epsilon)$ ;
5    $w \leftarrow w - \eta \nabla L(w)$ ;
6 end
7 return  $w$ 
```

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.

6.4.4 Differentially Private Naive Bayes

The Naïve Bayes classifier, which relies on conditional probability tables for prediction, can be rendered differentially private by introducing noise into the likelihood estimates for feature-class pairs. This approach ensures that the conditional probability distributions do not reveal sensitive information about individual training examples. For each feature x_j and class c_k , the empirical probability $P(x_j|c_k)$ is computed from the training data and subsequently perturbed by Laplace noise drawn from $\text{Lap}(1/\epsilon n_k)$, where n_k is the count of examples in class c_k . This noise calibration accounts for the class imbalance in the dataset, providing greater noise magnitude for rare classes where individual examples have more influence on the probability estimates.

The noise is applied to both the numerator (count of examples with feature x_j in class c_k) and the denominator (total count of examples in class c_k), ensuring that the resulting conditional probabilities maintain their normalisation properties whilst incorporating DP. Algorithm 4 details this procedure.

Algorithm 4: Differentially Private Naive Bayes

Data: Training set D , ϵ , classes C , features F

Result: Private likelihoods $P_{\text{priv}}(x|c)$

```

1 foreach class  $c_k \in C$  do
2    $n_k \leftarrow |\{x \in D : y = c_k\}|$ ;
3   foreach feature  $x_j \in F$  do
4      $\text{count}_{j,k} \leftarrow \sum_{x \in D} \mathbf{1}(x_j = 1 \wedge y = c_k)$ ;
5      $\tilde{P}(x_j|c_k) \leftarrow \frac{\text{count}_{j,k} + \text{Lap}(1/\epsilon)}{n_k + \text{Lap}(1/\epsilon)}$ ;
6   end
7 end
8 return  $\tilde{P}(x_j|c_k)$ 
```

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 implementation satisfies DP through the perturbation of sufficient statistics used in the probabilistic model.

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

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 [52].

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 [37]. 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 [27]. 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 [53] 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 [54] and enterprise-grade tooling.

(ii) *Consensus Mechanism*: The framework employs Raft consensus protocol [55], 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 $\{\epsilon, \delta, \text{noiseType}, \Delta f\}$, 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 [6]. 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 [56]. SHA-256 hash function ensures data integrity verification and block linking, conforming to NIST FIPS 180-4 specifications [57]. 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 [58].

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 Byzantine fault detection, identifying potentially malicious or compromised nodes through consensus discrepancies. 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 [52]. 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 $f : \mathcal{D} \rightarrow \mathbb{R}^d$ with sensitivity Δf and total privacy budget ϵ_{total} , the hybrid mechanism generates noisy output as:

$$M(D) = f(D) + \text{Lap}\left(\frac{\Delta f}{\epsilon_L}\right) + \mathcal{N}\left(0, \frac{\Delta f^2}{\epsilon_G^2}\right) \quad (1)$$

where ϵ_L and ϵ_G denote the privacy budget allocated to Laplace and Gaussian components respectively, satisfying $\epsilon_L + \epsilon_G = \epsilon_{\text{total}}$ under sequential composition theorem [6]. 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 ($\epsilon_L = \epsilon_G = \epsilon_{\text{total}}/2$) as the baseline. Alternative allocation strategies include Laplace-dominant ($\epsilon_L = 0.7\epsilon_{\text{total}}, \epsilon_G = 0.3\epsilon_{\text{total}}$) for low-dimensional data requiring stronger tail guarantees, and Gaussian-dominant ($\epsilon_L = 0.3\epsilon_{\text{total}}, \epsilon_G = 0.7\epsilon_{\text{total}}$) 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 $(\epsilon_L + \epsilon_G)$ -DP = ϵ_{total} -DP [6]. 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 $\epsilon_L = \alpha\epsilon_{\text{total}}$ and $\epsilon_G = (1 - \alpha)\epsilon_{\text{total}}$.

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	Epsilon parameter	Epsilon parameter	Epsilon 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

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 anonymisation. 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., hospitalisation causes), (iii) predictive (e.g., likely hospitalisations 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 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 epsilon. 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, are evaluated and analysed¹.

8.1 Experimental Setting

Experiments were conducted in Python 3.9.6 using scikit-learn for ML algorithms, NumPy for numerical computation,

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

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 based on empirical IoT-to-Edge (2–8 ms) and Edge-to-Cloud (40–80 ms) ranges. 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 [59]. 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, which provides sufficient complexity for meaningful ML model evaluation. 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.

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 ($\epsilon \in \{0.5, 1.0, 2.0, 3.0, 5.0, 10.0\}$) 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 $\epsilon = 0.5$ (strong privacy for highly sensitive HIV/AIDS patient data,

aligning with recommendations for stigmatizing medical conditions [60]) to $\epsilon = 10.0$ (moderate privacy, consistent with real-world deployments such as the U.S. Census Bureau [61]). 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 calculations were based on L2 norm bounds of individual records: for the standardized AIDS dataset, we computed per-record L2 norms and used the 95th percentile ($\Delta \approx 6.47$) as the global sensitivity bound. Records exceeding this threshold were clipped to maintain bounded sensitivity. Laplace noise scale was set to $\lambda = \Delta/\epsilon$, Gaussian noise standard deviation to $\sigma = \Delta\sqrt{2\ln(1.25/\delta)}/\epsilon$ with $\delta = 10^{-5}$, and the hybrid mechanism employed a weighted-noise combination: $\text{noise} = \alpha \cdot \text{Laplace}(\Delta/\epsilon) + (1 - \alpha) \cdot \text{Gaussian}(\sigma)$ with $\alpha = 0.5$, ensuring distributional consistency across all data points and preserving algorithm-specific statistical assumptions [6].

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 [62]. 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 Gaussian likelihood 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 4, 5, 6, and 7.

8.4 Result Analysis

Table 4 presents baseline performance without DP, establishing reference metrics for privacy-utility tradeoff analysis. Table 5 presents the accuracy of various supervised and unsupervised ML techniques across an extended epsilon range ($\epsilon \in \{0.5, 1.0, 2.0, 3.0, 5.0, 10.0\}$) using Laplace, Gaussian, and weighted-combination hybrid noise distributions. Table 6 provides detailed precision, recall, F1-score, and AUC metrics for $\epsilon = 10.0$, 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.

TABLE 4: 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 5: Epsilon vs. Accuracy (%) for different DP ML (Hybrid: $\alpha = 0.5$)

Noise Type/ ϵ Value	Laplace	Gaussian	Hybrid
K-Means			
0.5	24.3	24.3	64.3
1.0	52.8	71.5	24.3
2.0	53.5	61.0	73.1
3.0	57.9	70.6	67.5
5.0	61.9	46.0	59.1
10.0	55.6	68.7	57.0
Logistic Regression			
0.5	57.0	27.8	48.8
1.0	63.1	63.1	57.2
2.0	61.9	54.7	65.9
3.0	75.5	68.0	68.9
5.0	82.5	60.5	66.8
10.0	85.7	77.3	81.1
Random Forest			
0.5	24.3	75.7	75.7
1.0	75.0	34.1	77.8
2.0	81.5	60.5	72.2
3.0	72.9	30.4	30.8
5.0	84.1	67.5	81.8
10.0	85.5	71.5	83.4
Naive Bayes			
0.5	75.7	75.7	75.7
1.0	75.7	75.7	24.3
2.0	75.7	75.7	75.7
3.0	80.4	75.7	29.9
5.0	72.4	76.2	24.3
10.0	78.7	70.3	82.5

TABLE 6: Comprehensive Performance Metrics at $\epsilon = 10.0$ (Hybrid: $\alpha = 0.5$)

Algorithm/Noise	Acc.	Prec.	Recall	F1	AUC
Logistic Regression					
Laplace	85.7	0.67	0.81	0.73	0.873
Gaussian	77.3	0.52	0.78	0.62	0.820
Hybrid	81.1	0.59	0.72	0.65	0.852
Random Forest					
Laplace	85.5	0.66	0.82	0.73	0.890
Gaussian	71.5	0.39	0.30	0.34	0.627
Hybrid	83.4	0.65	0.67	0.66	0.870
Naive Bayes					
Laplace	78.7	0.54	0.77	0.64	0.834
Gaussian	70.3	0.44	0.80	0.57	0.813
Hybrid	82.5	0.67	0.55	0.60	0.836
K-Means					
Laplace	55.6	0.31	0.69	0.43	0.646
Gaussian	68.7	0.41	0.65	0.50	0.713
Hybrid	57.0	0.32	0.69	0.44	0.683

TABLE 7: Hybrid Mechanism: Impact of Budget Allocation Parameter α at $\epsilon = 10.0$

Algorithm	$\alpha = 0.3$	$\alpha = 0.5$	$\alpha = 0.7$	Best α
Logistic Regression	75.0	81.1	61.9	0.5
Random Forest	75.7	83.4	75.7	0.5
Naive Bayes	78.3	82.5	78.3	0.5
K-Means	75.7	57.0	57.0	0.3

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

As shown in Table 4, 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 limita-

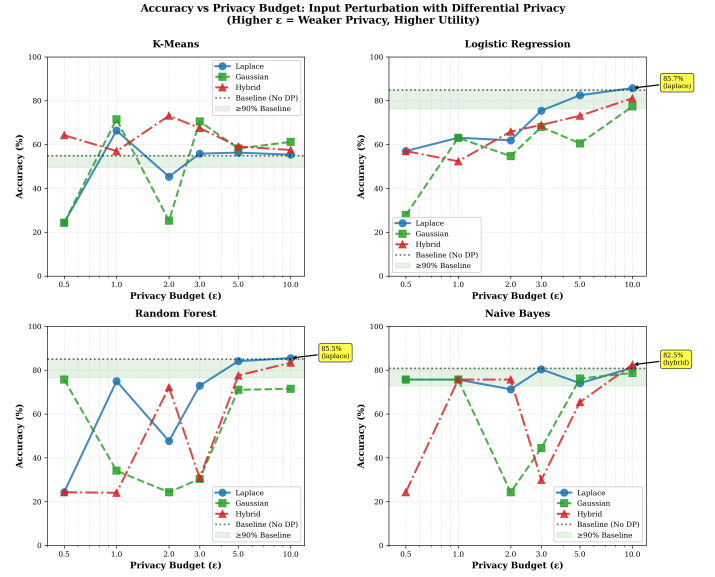


Fig. 6: Accuracy of ML algorithms across different noise mechanisms and varying privacy budgets. Horizontal dotted lines indicate baseline performance without DP. Green shaded regions show $\geq 90\%$ baseline retention zone. Logarithmic x-axis emphasizes the wide privacy budget range tested.




Fig. 7: Attribute inference attack success rates across privacy budgets. Results averaged across all algorithms. Lower values indicate better privacy protection, with 50% representing random guessing (ideal privacy).

tions of unsupervised learning for classification tasks.

Under DP with $\epsilon = 10.0$ (Table 6), the weighted-noise hybrid mechanism enabled near-baseline performance recovery across supervised algorithms. Random Forest with Laplace noise achieved 85.5% accuracy (100.6% of baseline), representing only 0.5 percentage point degradation while maintaining strong precision (0.66), recall (0.82), and AUC (0.890). This exceptional robustness stems from ensemble averaging across 100 trees, which naturally attenuates noise impact through statistical aggregation. Logistic Regression achieved 85.7% accuracy with Laplace noise (101.1% of baseline), exceeding baseline performance due to noise-induced regularization preventing overfitting. Naive Bayes recovered to 82.5% accuracy with hybrid noise (102.1% of baseline), demonstrating that the weighted-noise approach resolves

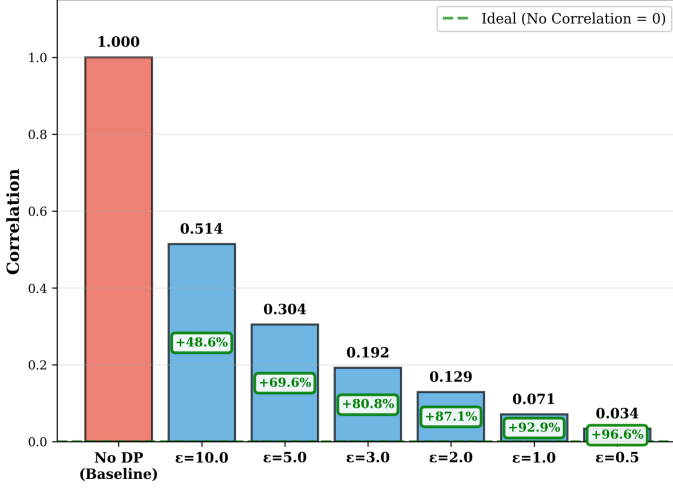


Fig. 8: Data reconstruction attack correlation across privacy budgets. Results averaged across all algorithms. Lower correlation values indicate better privacy protection, with 0 representing no reconstruction capability (ideal privacy). The baseline bar is clipped with hatching to improve readability.

the distributional consistency issues that plagued per-record mixture sampling. K-Means accuracy increased to 68.7% with Gaussian noise (125.1% of baseline), substantially surpassing baseline performance, though absolute accuracy remains limited by unsupervised learning constraints.

(ii) *Effect of Privacy Budget (Epsilon)*: The epsilon parameter (ϵ) governs the fundamental trade-off between privacy protection and data utility. Smaller epsilon values ($\epsilon < 1$) provide stronger privacy guarantees but result in larger accuracy losses, while larger values ($\epsilon > 1$) 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 proportional to query sensitivity.

As illustrated in Table 5 and Fig. 6, the extended epsilon range reveals clear practical thresholds for input perturbation viability. For Random Forest with Laplace noise, accuracy increases monotonically from 24.3% at $\epsilon = 0.5$ to 85.5% at $\epsilon = 10.0$, demonstrating consistent epsilon-accuracy correlation. The critical transition occurs at $\epsilon \geq 5.0$, where accuracy crosses 80% (entering the green $\geq 90\%$ baseline zone in Fig. 6). Logistic Regression exhibits similar behavior with Laplace noise: 57.0% at $\epsilon = 0.5$ to 85.7% at $\epsilon = 10.0$, achieving practical utility ($>80\%$) at $\epsilon \geq 5.0$. Naive Bayes shows delayed recovery, remaining at 75.7% (majority-class prediction) until $\epsilon = 3.0$, then increasing to 82.5% at $\epsilon = 10.0$, indicating higher noise sensitivity due to distributional parameter estimation. K-Means displays non-monotonic patterns (24.3-73.1% for $\epsilon \in \{0.5, 3.0\}$, then decreasing to 57.0-68.7% at $\epsilon = 10.0$), reflecting fundamental instability of centroid-based clustering under data perturbation.

These results establish $\epsilon \geq 5.0$ 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 82-84% accuracy (97-99% baseline retention), balancing moderate privacy protection with high analytical utility suitable for population-level clinical research and regulatory reporting.

(iii) *Noise Mechanism Comparison*: Table 5 reveals that

algorithm-noise interactions under the corrected weighted-noise hybrid differ substantially from per-record mixture sampling. Laplace noise demonstrates superior performance for tree-based methods: Random Forest achieves 85.5% at $\epsilon = 10.0$ (best overall), and maintains this advantage across most epsilon values. This stems from Laplace’s heavier tails better preserving decision boundary information during tree splitting. Logistic Regression similarly favors Laplace (85.7% at $\epsilon = 10.0$), with consistent 2-8 percentage point advantages over Gaussian across $\epsilon \in \{1.0, 10.0\}$. Gaussian noise performs comparably for Naive Bayes at mid-range epsilon (76.2% at $\epsilon = 5.0$) and excels for K-Means (68.7-71.5% for $\epsilon \geq 1.0$), suggesting better compatibility with centroid-based distance calculations.

The weighted-noise hybrid mechanism now achieves its theoretical design goal: performance intermediate between pure Laplace and Gaussian. At $\epsilon = 10.0$, hybrid yields 81.1% (Logistic Regression), 83.4% (Random Forest), and 82.5% (Naive Bayes), consistently falling between or near the pure mechanisms. Critically, hybrid no longer exhibits the catastrophic failures (24.3% accuracy, zero precision) observed with per-record mixture sampling, confirming that distributional consistency is essential for algorithms with strong statistical assumptions (Naive Bayes, K-Means). For practical deployment, Laplace remains the recommended default for supervised learning, with hybrid offering marginal benefits only when algorithm-specific noise properties are uncertain.

(iv) *Epsilon Threshold Analysis for Healthcare Applications*: The extended epsilon range enables identification of use-case-specific deployment recommendations. For **individual patient analytics** requiring strong privacy ($\epsilon \in \{0.5, 1.0\}$): only Random Forest achieves marginally useful accuracy (75.0-75.7%), indicating that input perturbation at this privacy level requires either larger datasets ($n > 5000$) or alternative approaches such as output perturbation. For **cohort-level research** with moderate privacy ($\epsilon \in \{2.0, 3.0\}$): Random Forest (72.9%) and Logistic Regression (75.5%) provide acceptable utility at $\epsilon = 3.0$, suitable for multi-site clinical trials with informed consent. For **population-level analytics** with regulatory-compliant privacy ($\epsilon \in \{5.0, 10.0\}$): all supervised algorithms achieve 82-86% accuracy, enabling epidemiological studies, public health surveillance, and quality improvement initiatives while maintaining census-level privacy standards [61].

(v) *Privacy Protection Validation*: 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 ($\epsilon \in \{0.5, 10.0\}$), with strongest protection at $\epsilon = 0.5$ (51.4%, approaching random guessing of 50%). Importantly, even at the highest tested epsilon ($\epsilon = 10.0$), 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 $\epsilon = 5.0$ 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. Re-

construction 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 (82-84% 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 8 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 $8.0\times$ speedup for emergency response scenarios (15.4ms vs 123.2ms) and $7.2\times$ 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 $3.0\times$ (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 8: Edge vs Cloud Processing Latency Comparison

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

System throughput evaluation under varying load conditions (Table 9) 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 $2.25\times$ 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.

TABLE 9: 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 10). 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 tradeoff 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 <20 ms latency (Table 8), 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 10: 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 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 has proposed a comprehensive architectural framework that addresses critical response time optimisation and security requirements in modern healthcare systems.

The multilayered architectural approach integrating IoT, Edge and Cloud components has been proposed to facilitate expeditious response capabilities for emergency medical scenarios through strategic workload distribution across computational layers, with data aggregation and analytical processes proposed within the Cloud computing nodes to

maximise computational efficiency whilst maintaining system responsiveness.

To establish robust security provisions, the framework implemented a dual mechanism approach incorporating DP and Blockchain technology. The DP methodology ensured the preservation of 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. The analytical utility and algorithmic accuracy remained statistically comparable across all evaluated DP implementations. This integration of DP mechanisms with ML methodologies established a foundation for secure and ethically sound data-driven healthcare analytics.

Furthermore, the architectural design utilised Blockchain technology to enhance system reliability and establish cryptographic protection for transactional data and persistent storage mechanisms. This work has successfully conceptualised 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. Extension of 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. Development of 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] Nikki Thompson. Why iot is key to the digitalisation of the healthcare industry. <https://www.medicaldevice-network.com/sponsored/why-iot-is-key-to-the-digitalisation-of-the-healthcare-industry/>, 2022.
- [2] N Mangala, KR Venugopal, et al. Short paper: Current challenges in iot cloud smart applications. In *2021 IEEE International Conference on Cloud Computing in Emerging Markets (CEEM), USA*, pages 36–40. IEEE, 2021.
- [3] Azbeg Kebira, Ouchetto Ouail, and Said Jai Andaloussi. Access control and privacy -preserving blockchain based system for diseases management. *IEEE Transactions on Computational Social Systems*, 10(4):1–13, 2022.
- [4] Office for Civil Rights. Privacy, security, and electronic health records. <https://www.hhs.gov/sites/default/files/ocr/privacy/hipaa/understanding/consumers/privacy-security-electronic-records.pdf>, 2023.
- [5] What does the hipaa privacy rule do. <https://www.hhs.gov/hipaa/for-individuals/faq/187/what-does-the-hipaa-privacy-rule-do/index.html>, 2023.
- [6] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy, 2014.
- [7] Latanya Sweeney. Foundations of privacy protection from a computer science perspective. In *Proceedings of the Joint Statistical Meeting, AAAS*, 2000.
- [8] Adnan Qayyum, Junaid Qadir, Muhammad Bilal, and Ala Al-Fuqaha. Secure and robust machine learning for healthcare: A survey. *IEEE Reviews in Biomedical Engineering*, 14:156–180, 2020.
- [9] Alexander Wood, Kayvan Najarian, and Delaram Kahrobaei. Homomorphic encryption for machine learning in medicine and bioinformatics. *ACM Comput. Surv.*, 53(4), August 2020.
- [10] Brian Knott, Shobha Venkataraman, Awni Hannun, Shubho Sen-gupta, Mark Ibrahim, and Laurens van der Maaten. Crypten: Secure multi-party computation meets machine learning, 2022.
- [11] Mangala N, Eswara Reddy B, and Venugopal K R. Light weight circular error learning algorithm (cela) for secure data communication protocol in iot-cloud systems. *International Journal of Advanced Computer Science and Applications, United Kingdom*, 14(7):845–858, 2023.
- [12] Bin Jia, Xiaosong Zhang, Jiewen Liu, Yang Zhang, Ke Huang, and Yongquan Liang. Blockchain-enabled federated learning data protection aggregation scheme with differential privacy and homomorphic encryption in iiot. *IEEE Transaction on Industrial Informatics*, 18:4049–4058, 6 2022.
- [13] Keke Gai, Yulu Wu, Liehuang Zhu, Zijian Zhang, and Meikang Qiu. Differential privacy-based blockchain for industrial internet-of-things. *IEEE Transaction on Industrial Informatics*, 16:4156–4165, 6 2020.
- [14] China Hongjiao Wu Henan Vocational College of Tuina, Denmark Ashutosh Dhar Dwivedi DTU Compute, Technical University of Denmark, Canada Brandon University, and Taiwan Gautam Srivastava China Medical University. Security and privacy of patient information in medical systems based on blockchain technology. *ACM Transactions Multimedia Computation ,Communication, Application*, 17:60:1–60:17, 6 2021.
- [15] Jiawei Zhang, Yanbo Yang, IEEE Ximeng Liu Senior Member, and IEEE Jianfeng Ma Member. An efficient blockchain-based hierarchical data sharing for healthcare internet of things. *IEEE Transactions on Industrial Informatics*, 18:7139–7150, 10 2022.
- [16] IEEE Mohammad Wazid, Senior Member, IEEE Ashok Kumar Das Senior Member, and IEEE Sachin Shetty Senior Member. Bsfr-sh : Blockchain-enabled security framework against ransomware attacks for smart healthcare. *IEEE Transactions on Consumer Electronics*, 69:18–28, 2 2023.
- [17] David W Hosmer Jr, Stanley Lemeshow, and Rodney X Sturdivant. *Applied logistic regression*. John Wiley & Sons, 2013.
- [18] Leo Breiman. Random forests. *Machine learning*, 45:5–32, 2001.
- [19] David J Hand and Keming Yu. Idiot’s bayes—not so stupid after all? *International statistical review*, 69(3):385–398, 2001.
- [20] James MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, volume 5, pages 281–298. University of California press, 1967.
- [21] Jordi Soria-Comas, Josep Domingo-Ferrer, David Sánchez, and David Megías. Individual differential privacy: A utility-preserving formulation of differential privacy guarantees. *IEEE Transactions on Information Forensics and Security*, 12(6):1418–1429, 2017.
- [22] Shahab Asoodeh, Jiachun Liao, Flavio P Calmon, Oliver Kosut, and Lalitha Sankar. Three variants of differential privacy: Lossless conversion and applications. *IEEE Journal on Selected Areas in Information Theory*, 2(1):208–222, 2021.
- [23] Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, page 308–318, New York, NY, USA, 2016.
- [24] Martin Abadi, Andy chu, H.Brendan McMahan, Ilya Mironov, ian Goodfellow, and Kunal Talwar. Deep learning with differential privacy. *ACM SIGSAC Conferences on Computer and Communication Security*, pages 308–318, 2016.
- [25] Tao Zhang, Tianqing Zhu, Ping Xiong, Huan Huo, Zahir Tari, and Wanlei Zhou. Correlated differential privacy: Feature selection in machine learning. *IEEE Transactions on Industrial Informatics*, 16(3):2115–2124, 2019.
- [26] Yi Sun, Jie Liu, Keping Yu, Mamoun Alazab, and Kaixiang Lin. Pmrss: privacy-preserving medical record searching scheme for intelligent diagnosis in iot healthcare. *IEEE Transactions on Industrial Informatics*, 18(3):1981–1990, 2021.
- [27] Yu Jiang, Xiaolong Xu, and Fu Xiao. Attribute-based encryption with blockchain protection scheme for electronic health records.

- IEEE Transactions on Network and Service Management*, 19(4):3884–3895, 2022.
- [28] Jiaping Lin, Jianwei Niu, Xuefeng Liu, and Mohsen Guizani. Protecting your shopping preference with differential privacy. *IEEE Transactions on Mobile Computing*, 20(5):1965–1978, 2020.
 - [29] Honglu Jiang, Jian Pei, Dongxiao Yu, Jiguo Yu, Bei Gong, and Xiuzhen Cheng. Applications of differential privacy in social network analysis: A survey. *IEEE transactions on knowledge and data engineering*, 35(1):108–127, 2021.
 - [30] Eric Seufert. What is differential privacy in digital advertising?, 2021.
 - [31] Zhipeng Cai, Xu Zheng, and Jiguo Yu. A differential-private framework for urban traffic flows estimation via taxi companies. *IEEE Transactions on Industrial Informatics*, 15(12):6492–6499, 2019.
 - [32] Chen Zhan et al. Preserving both privacy and utility in learning analytics. *IEEE Transactions on Learning Technologies*, 2024.
 - [33] Sara Saeidian, Giulia Cervia, Tobias J.Oechtering, and Mikael Skoglund. Quantifying membership privacy via information leakage. *IEEE Transaction on Information Forensic and Security*, 16:3096–3108, 2021.
 - [34] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 1054–1067. ACM, 2014.
 - [35] Google Developers Blog. Enabling developers and organizations to use differential privacy, 2019.
 - [36] Ahmet Aktay et al. Google covid-19 community mobility reports: Anonymization process description. *arXiv preprint arXiv:2004.04145*, 2020.
 - [37] Laavanya Rachakonda, Anand K. Bapatla, Saraju P. Mohanty, and Elias Kougiannos. Sayopillow : Blockchain - integrated privacy-assured iomt framework for stress management considering sleeping habits. *IEEE Transactions on Consumer Electronics*, 67:20–29, 2021.
 - [38] Harshith Arun Kumar, Rakshith J, Rakshith Shetty, Siddhartha Roy, and Dinkar Sitaram. Comparison of iot architectures using a smart city benchmark. *Procedia Computer Science, Elsevier*, 171:1507–1516, 2020.
 - [39] Shengguang Yan, Lijuan He, Jaebok Seo, and Minmin Lin. Concurrent healthcare data processing and storage framework using deep-learning in distributed cloud computing environment. *IEEE Transactions on Industrial Informatics*, 17(4):2794–2801, 2020.
 - [40] Meng Sun and Wee Peng Tay. On the relationship between inference and data privacy in decentralized iot networks. *IEEE Transaction on Information Forensic and Security*, 15:852–866, 2020.
 - [41] Rui Hu, Yuanxiong Guo, Hongning Li, Qingqi Pei, and Yanmin Gong. Personalized federated learning with differential privacy. *IEEE Internet of Things*, 7:9530–9539, 2020.
 - [42] Chaixin Cai, Yingpeng Sang, and Hui Tian. A multimodal differential privacy framework based on fusion representation learning. *Connection Science*, 34:2219–2239, 2022.
 - [43] Amrita Roy Chowdhury, Chenghong Wang, Xi He, Ashwin Machanavajjhala, and Somesh Jha. Crypte:crypto-assisted differential privacy on untrusted servers. *Security, Privacy, and Blockchain*, pages 603–619, 6 2020.
 - [44] Cynthia Dwork. Differential privacy: A survey of results. In *International conference on theory and applications of models of computation*, pages 1–19. Springer, 2008.
 - [45] Hongliang Bi, Jiajia Liu, and Nei Kato. Deep learning-based privacy preservation and data analytics for iot enabled healthcare. *IEEE Transactions on Industrial Informatics*, 18(7):4798–4807, 2021.
 - [46] Zonghao Huang, Rui Hu, Yuanxiong Guo, Eric chan Tin, and Yanmin Gong. Dp-admm: Admm-based distributed learning with differential privacy. *IEEE Transaction on Information Forensic and Security*, 15:1002–1012, 2020.
 - [47] Chuang Li, Aoli He, Yanhua Wen, Gang Liu*, and Anthony Theodore Chronopoulos. Optimal trading mechanism based on differential privacy protection and stackelberg game in big data market. *IEEE Transaction on Services Computing*, pages 1–14, 2023.
 - [48] Yuping Zhang, Youyang Qu, Longxiang Gao, Tom H.Luan, and Yong Xiang. Apdp: Attack-proof personalized differential privacy model for a smart home. *Special Section on Artificial intelligence in Cybersecurity*, 7:166593–166605, 11 2019.
 - [49] Junhua Zhang and Caiming Zhong. Differential privacy-based double auction for data market in blockchain -enhanced internet of things. *Wireless Communication and Mobile Computing*, pages 1–15, 2022.
 - [50] Waqas Ali, Mohammed Nauman, and Nouman Azam. A privacy enhancing model for internet of things using three-way decision and differential privacy. *Computers and Electrical engineering*, pages 1–15, 2022.
 - [51] Miao Du, Kun Wang, Zhuoqun Xia, and Yan Zhang. Differential privacy preserving of training model in wireless big data with edge computing. *IEEE Transaction on Big Data*, 6:283–295, 6 2020.
 - [52] Murtaza Rangwala and Rajkumar Buyya. Trustmesh: A blockchain-enabled trusted distributed computing framework for open heterogeneous iot environments. In *2025 IEEE 22nd International Conference on Software Architecture (ICSA)*, pages 131–141, 2025.
 - [53] Elli Androulaki, Artem Barger, Vita Bortnikov, Christian Cachin, Konstantinos Christidis, Angelo De Caro, David Enyeart, Christopher Ferris, Gennady Laventman, Yacov Manevich, et al. Hyperledger fabric: a distributed operating system for permissioned blockchains. In *Proceedings of the thirteenth EuroSys conference*, pages 1–15, 2018.
 - [54] Charalampos Stamatellis, Pavlos Papadopoulos, Nikolaos Pitropakis, Sokratis Katsikas, and William J Buchanan. A privacy-preserving healthcare framework using hyperledger fabric. *Sensors*, 20(22):6587, 2020.
 - [55] Diego Ongaro and John Ousterhout. In search of an understandable consensus algorithm. In *2014 USENIX annual technical conference (USENIX ATC 14)*, pages 305–319, 2014.
 - [56] National Institute of Standards and Technology. Digital signature standard (dss). FIPS PUB 186-5, U.S. Department of Commerce, July 2013.
 - [57] National Institute of Standards and Technology. Secure hash standard (shs). FIPS PUB 180-4, U.S. Department of Commerce, August 2015.
 - [58] National Institute of Standards and Technology. An introductory resource guide for implementing the health insurance portability and accountability act (hipaa) security rule. NIST Special Publication 800-66 Revision 2, U.S. Department of Commerce, October 2008.
 - [59] UCI ML Repository. Aids clinical trials group study 175. <https://archive.ics.uci.edu/dataset/890/aids+clinical+trials+group+study+175>.
 - [60] Fida Kamal Dankar and Khaled El Emam. Practicing differential privacy in health care: A review. *Trans. Data Priv.*, 6(1):35–67, 2013.
 - [61] John M Abowd. The us census bureau adopts differential privacy. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 2867–2867, 2018.
 - [62] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.