



Evaluating Privacy Leakage Risks in High Capacity Healthcare Analytics Models

Safikur Rahman Khan¹, Mohd Waris Khan^{2*}, Asif Khan³

¹Department of Computer Application, Integral University, Dasauli, Bas-ha Kursi Road, Lucknow, Uttar Pradesh, 226026, India.
E-mail: srkhan@iul.ac.in

²Department of Computer Application, Integral University, Dasauli, Bas-ha Kursi Road, Lucknow, Uttar Pradesh, 226026, India.
E-mail: mwariskhan@iul.ac.in

³Department of Computer Application, Integral University, Dasauli, Bas-ha Kursi Road, Lucknow, Uttar Pradesh, 226026, India.
E-mail: asifk@iul.ac.in

Abstract

The fast development of digital healthcare systems and electronic health records has created opportunities for improved decision making based on data that was previously inaccessible. However, the sensitive nature of health information raises important questions regarding security, privacy and ethical compliance. This paper seeks to introduce an optimized machine learning model for health data analytics that is privacy-aware. The approach uses techniques such as differential privacy, federated learning and safe aggregation for collaborative health data analysis without disclosing raw patient data. The proposed model manages the privacy-utility trade-off while lowering the communication cost via gradient clipping, noise injection and effective optimization methods. An experimental evaluation on benchmark healthcare datasets shows that the model provides a high level of privacy guarantees and a high degree of prediction accuracy. This is quantified in terms of differential privacy parameters (ϵ , δ). The system also improves resilience against adversarial attacks and scalability, thus providing reliable results suitable for clinical decision support. Our study contributes toward the creation of a secure, effective and moral solution for the next generation of health data analytics by combining optimization algorithms with privacy-preserving safeguards.

Keywords: Health Data Analytics, Privacy-Preserving Machine Learning, Differential Privacy, Secure Aggregation.

1. Introduction

The rapid growth of healthcare data, powered by advances in Electronic Health Records (EHRs), where sensors, medical imaging, and genetic technologies, has generated significant prospects in Machine Learning (ML) to improve disease diagnosis, treatment methodologies, and precision medicine. The sensitive and personally identifiable characteristics of medical data engender significant concerns over privacy, security and adherence to regulations such as HIPAA and GDPR [1]. To maintain ethical accountability and public trust, ML in healthcare must be executed with resilient Privacy-Preserving algorithms. In response to this critical necessity, Privacy-Preserving Machine Learning (PPML) has evolved as a viable method for safeguarding patient data while preserving model performance. A paradigm of the fourth industrial revolution, the demand for technologies that process and analyse big data in the healthcare sector has led to active research in the field [2]. However, as the use of medical data increases, issues regarding the protection of personal information in the data are being raised. The most prominent issue is the area of personal information protection, as electronic records and billing data used in research contain personal information, such as patient gender, age and address. This has led to many medical data breaches, and cases of personal information breaches caused by this can be found in many papers. Many countries have enacted laws to prevent such information from leaking, such as the Personal Information Protection Act (PIPA) in Korea and the General Data Protection Regulation (GDPR) in Europe.

Due to such laws, conducting epidemiological research using patients' personal information contradicts protecting their personal information. Therefore, various attempts are being made to analyse medical data while protecting patients' personal information, specifically using privacy protection technologies [3, 4]. Examples of such technologies are differential privacy (DP) and homomorphic encryption. More recently, concepts, such as blockchain and federated learning, have also been utilized to ensure the privacy of personal information. When blockchain is decentralized, it shows excellent performance in the security sector, indicating that patients' data can be used safely. Blockchain has also been applied in the field of drug supply network management. However, since the risk of counterfeiting is high, it is applied only for certification, sales and distribution monitoring. For the initial application, drug transactions were monitored by storing all the information on drug transactions in a blockchain. Another technology, federated learning, enables data diversity when applied in the medical field. Federated learning improves

artificial intelligence (AI) learning for patients with similar patient symptoms by connecting data from multiple medical institutions. The more data collected for similar patients, the greater the predictive accuracy of AI for the patient group [5]. This can provide valuable insight to clinicians along with individual physician findings.

The remainder of this paper is organized as follows: Section 2 reviews the relevant literature on healthcare data analytics, machine learning, and related technologies. Section 3 discusses ethical considerations and privacy-oriented approaches for secure health data analysis. Section 4 presents the proposed methodology, including data pre-processing, model development, and evaluation procedures. Section 5 presents and analyses the experimental results obtained from the proposed approach. Finally, Section 6 discusses the key findings, limitations and implications of the study and concludes the paper with potential directions for future research.

2. Review of literature

The current body of research on health data analytics underscores the vital roles of machine learning (ML), data privacy, and healthcare informatics. Traditional ML models are recognized for their effectiveness in illness prediction, clinical decision-making, and patient monitoring by leveraging electronic health records and medical databases. However, there are significant concerns around data security and privacy that necessitate adherence to regulatory frameworks like HIPAA and GDPR. In response, federated learning has emerged as a decentralized method that permits collaboration among healthcare institutions to train models while safeguarding sensitive patient information. This is complemented by the concept of differential privacy, which introduces statistical noise to model updates, thus enhancing privacy. Advances in safe multi-party computations and cryptographic algorithms are also key to securing medical data during collaborative analyses. Research by Aminizadeh et al. (2024) highlights the potential of artificial intelligence and distributed systems to transform healthcare services [6]. Yet, challenges pertaining to scalability, interoperability, and privacy remain. Their research paves the way for the incorporation of sophisticated AI systems into practical healthcare applications. Meanwhile, Cusi et al. (2022) have established thorough clinical guidelines for the diagnosis and management of non-alcoholic fatty liver disease, showcasing the critical need for expert consensus to standardize approaches across general care and endocrinology, thereby positively influencing patient outcomes [7]. Moreover, Masoumian Hosseini et al. (2023), emphasize the integration of telemedicine within digital innovation, arguing for a patient-centered design that merges telemedicine effectively with traditional healthcare frameworks, thereby ensuring both accessibility and continuity of care [8]. Their insights into wearable technologies, especially smartwatches, illustrate their emerging significance in healthcare by facilitating continuous monitoring and support. These consumer-oriented devices are progressively morphing into medical instruments that advance preventive care, chronic disease management, and real-time health tracking. Ye et al. (2021) advocate for integrated learning-based systems to handle large volumes of medical data, assessing integrated algorithms against the complexity of health statistics to improve the precision of analytics and clinical decision-making [9]. Additionally, Rehman et al. (2022) present an expansive overview of big data analysis in healthcare, exploring prevailing trends, future directions, and the challenges associated with processing vast amounts of sensitive data, while addressing privacy concerns and infrastructural limitations [10]. Khanna et al. (2023) demonstrate innovative uses of the Internet of Things combined with deep learning for enhancing disease detection through biological ECG data, aiming to enhance diagnostic accuracy and efficiency in cardiology [11]. Meanwhile, Elshewey et al. (2024) focus on developing optimized algorithms for plant disease detection using deep learning, which, despite their agricultural focus, draw parallels to healthcare by linking food security, agricultural sustainability, and human health [12]. In conclusion, the findings from 2021 to 2024 illuminate emerging trends in healthcare innovation, marked by the synthesis of artificial intelligence, wearable technology, telemedicine, big data analytics, the Internet of Things, and sophisticated algorithms, all poised to shape the future of healthcare. Nevertheless, these advancements bring forth critical challenges regarding privacy, interoperability, and ethical integration, stressing the ongoing need for research focused on secure and efficient health data analytics models.

3. Ethical and Privacy-Oriented Health Data Analysis

Massive digital health ecosystems have emerged, creating new possibilities for personalized care, predictive modeling and clinical innovation. But these advances also provide a crucial challenge to ensure that health data are gathered and handled in a manner that respects privacy and is ethically sound. Ethical and privacy-oriented health data analysis is a framework for using sensitive medical data that combines technical safeguards with adherence to moral principles such as respect for autonomy, fairness, transparency and accountability. In this regard, the protection of patient privacy is not just a technical issue but also a moral obligation that sustains the trust that exists between patients, researchers and healthcare providers [9]. One of the most critical ethical challenges in health data analytics is informed consent. Patients need to be told what will happen to their data, who will do it, and what the purpose will be. Ethical data practices

include voluntary participation and the ability for people to delete their data if they want to. However, it is difficult to gain meaningful consent in modern healthcare when data are typically integrated from several sources such as wearable sensors, telemedicine platforms, electronic health records and genetic information. New privacy-protecting solutions, such as federated learning, allow institutions to train machine learning algorithms together without centralising patient information and mitigate ethical problems connected to permission violations and data exploitation [10]. Another important topic is the role of privacy enhancing technologies (PETs). Methods such as differential privacy add carefully measured noise to the input or the output of the model, making it impossible to reverse engineer information on an individual level. Homomorphic encryption and safe multi-party computation give cryptographic guarantees that allow many parties to work on encrypted data without leaking any of the underlying values [8, 13]. These technologies maintain ethical values by providing anonymity and beneficial information for medical research and practice. Both fairness and bias avoidance are essential from the ethical perspective. If the underlying datasets are biased or unequal, machine learning algorithms trained on health data may inadvertently repeat or amplify existing discrepancies in treatment. To prevent patients from receiving different treatments because of their age, gender, race or socioeconomic situation, biases in ethical analysis must be detected and corrected [7, 11]. Enhanced equity in health data analytics using fairness-aware learning algorithms and public reporting of model performance across demographic groups. Regulatory systems such as GDPR in Europe, HIPAA in the US, and new international health data governance standards provide the legal boundaries within which patient privacy must be protected. But ethical and privacy-conscious health data analysis is more than compliance, and requires continuing surveillance, auditing and accountability procedures for sustained dependability. Robust governance frameworks for healthcare companies that document data ownership, create data sharing agreements, and have a process for redress if there is misuse or breaches, are critical. Data analysis should significantly improve health outcomes without violating human rights, consistent with the principle of beneficence, which seeks to maximize benefits and minimize harm. For example, in using predictive analytics in an attempt to detect early signals of chronic disease, it is critical to ensure that the models are accessible and useful for healthcare practitioners, but also shield patients from unnecessary risks of stigmatization or data breaches [6]. The consensus is that explainable AI approaches are a trade-off between practical utility and ethical transparency, seeking to better understand how patients and doctors make decisions. Sustainability and future readiness are the priorities for ethical and privacy-centric health data analysis. The growing volume of multimodal data generated by health care systems necessitates scalable approaches that preserve privacy and allow for study. There is a need for ethical foresight to anticipate challenges such as secondary use of data, cross-border flows of data, and the introduction of new technologies such as genomics and AI-based diagnostics. Embedding ethical reasoning into the design of privacy-preserving algorithms and governance structures [12] can help health systems ensure that technological development is aligned to the fundamental values of patient-centred care and human dignity with fair access and public trust as the end goal. A holistic strategy that brings together legal compliance, state-of-the-art privacy technology, fairness-aware AI, transparent governance and ethical responsibility is required to conduct an ethical and privacy-aware analysis of health data. In this way, healthcare analytics will be both innovative and trustworthy, protecting the individual while serving the good of society as a whole.

4. Methodology

The proposed privacy-preserving framework for the health data analytics is designed as a systematic multi-phase process (Figure 1) considering both technological and ethical aspects to provide safe and effective outcomes. The first phase of the process is data collection and pre-processing. Data are collected from multiple health data sources such as electronic health records (EHRs), lab test outcomes, imaging datasets, and data generated from Internet of Things (IoT) devices such as wearable sensors and smartwatches to mimic real-world healthcare environments. The preprocessing steps include removing redundant or irrelevant data, encoding categorical features into machine-readable formats, normalizing and standardizing attributes to facilitate consistency, and statistically imputing missing values [14, 15]. This phase ensures the conversion of the unstructured health data to high quality structured data set used in machine learning. The second phase of model creation mainly uses Federated learning (FL). Federated learning allows each participating healthcare facility to train models locally on their own data rather than aggregating sensitive patient data into a single repository. Instead of raw data, model parameters are securely sent to a central aggregator. This decentralized architecture enables collaborative learning on many datasets while preserving data secrecy. The model architecture is chosen according to the kind of health data: ensemble learning techniques for structured tabular data, recurrent neural networks for time-series wearable data, and deep neural networks for imaging. The learning method incorporates privacy-preserving techniques in the third step. Differential privacy is achieved by adding controlled Gaussian noise and clipping gradients to a pre-defined threshold before the updates are communicated to the global

aggregator. This ensures that the outputs of the model cannot be used to reconstruct the individual-level information.




Figure 1: Methodology

Moreover, secure aggregation techniques are employed to encrypt the updates in transmission, so that the server can only access the aggregated results, not the individual contributions. These measures together provide strong guaranties for the mathematical privacy, while preserving the utility of the algorithm. In the fourth stage, optimization techniques are used to improve performance. During the federated training, the computing load and communication overhead are minimized through gradient compression, adaptive learning rates and early halting. The strong aggregation strategies such as median-based aggregation are used to protect against malicious or inaccurate client updates. The total privacy budget is monitored by a privacy accountant to ensure that the training process maintains the proper balance between privacy protection and model accuracy. The proposed framework is evaluated and validated in the fifth stage. The prediction accuracy, sensitivity, specificity, F1-score and area under the ROC curve (AUC) of the optimized model are evaluated using benchmark healthcare datasets. System-level performance is evaluated in terms of scalability, computational efficiency, and communication cost, while privacy effectiveness is evaluated by differential privacy parameters. We perform subgroup analyzes to detect any algorithmic bias based on demographic factors such as age, sex and ethnicity for fairness and ethical concerns. We also conduct adversarial simulations such as membership inference and model inversion attacks to assess the robustness of the framework against potential privacy risks. In the sixth stage, we compare baseline models, including centralized machine learning without privacy restrictions and federated models without privacy mechanisms. The comparison demonstrates the advantages of the proposed framework in achieving a trade-off between strong privacy preservation and high predictive utility. The method ends with a discussion of the implications of deploying such models in real-world clinical settings, including scalability, regulatory compliance and the long-term sustainability of ethical, privacy-preserving health data analytics.

4.1 Data Description

This synthetic healthcare dataset is meticulously built with help of Python's Faker package, to resemble the qualities of genuine patient records, with perfect privacy and ethical security [15]. The main purpose of this dataset is to provide data science, machine learning and healthcare analytics learners, researchers and practitioners a realistic, but safe, environment to experiment and build models. This synthetic dataset is especially useful since genuine healthcare data is sensitive and frequently unavailable owing to privacy laws and regulations such as HIPAA and GDPR. The dataset maintains the structure and variety of clinical data but does not include information that may disclose real patient information. The dataset comprises patient-level and hospital-level records and is therefore equipped to support a broad variety of analytical and predictive applications. Each row of the dataset corresponds to one patient admission, while the columns describe several facets of the healthcare experience including demographics, diagnosis, treatment and results. Patient identities, names, genders and ages were synthetically produced to reflect actual demographic variety. Contact information (addresses and phone numbers) have also been included to offer some idea of geographical distribution which may be useful for regional or spatial health analysis. Admission data like admission ID, admission date and discharge date provide a calendar of hospital visits, allowing users to assess hospital stay lengths and patient flow. The departments and kinds of wards have been provided so as to show the varied levels of care for example regular wards, private rooms and critical care units. The medical diseases and diagnoses have been created to mirror a spectrum of clinical circumstances and the lab test results and diagnostic reports mimic quantifiable health markers making the data set extremely useful for classification, regression and risk prediction activities. The treatment and procedure fields capture synthetic but realistic healthcare treatments that might be associated with outcomes such as recovery, referral or death.

Information on physicians or consultants is also supplied to enable research on doctors' workload and efficacy of treatments. From the financial side, billing amounts and insurance status variables depict the economic part of healthcare providing users opportunities to analyse costs, insurance claims and payment mechanisms. The columns for outcome and discharge status allow the evaluation of the treatment outcomes and hence make the dataset useful for supervised machine learning tasks to predict patient recovery or identify risk factors. This mix of traits results in a dataset that is rich in information and adaptable in use. Users may use it for exploratory data analysis, healthcare operations research, building standard ML models, or even for privacy preserving AI projects like federated learning and differential privacy. This synthetic dataset combines demographic, clinical, administrative, and financial components, reflecting the multifaceted characteristics of real-world healthcare data, while preserving its educational and ethical significance.

5. Experimental Result

In this experimental investigation, we executed the operations on the selected dataset using three different ML models, namely k-Nearest Neighbors (kNN), Neural Network and Decision Tree. The aim of using different models was to check what results we will get by applying distance-based learning, deep learning and rule-based learning to the same problem. Each of these models has a different approach to pattern recognition, kNN use similarity metrics in the feature space, Neural Networks use nonlinear transformations with hidden layers and Neural Networks use a hierarchically structured set of decision rules based on feature splits. The testing of these three different models allowed us to provide a comparison assessment that contributes to the depth and trustworthiness of our technique. For model performance assessment, we used two crucial evaluation measures which are accuracy and F1-score. Accuracy is a simple metric to quantify how many predictions were categorised correctly out of the total samples. But while accuracy alone may often be deceptive (particularly with unbalanced datasets), we also computed the F1-score, which incorporates precision and recall. Inclusion of the F1-score led to a more fair evaluation of the models' efficacy, especially in relation to their capacity to detect both the majority and the minority classes. The use of dual-metric assessment was important for robustness of our findings. The findings indicated considerable disparities among the models. The kNN model was used as a baseline. It is pretty basic and restricted in terms of performance since it relies significantly on the choice of parameter k and distance metric, but it gave a valuable point of reference to compare with. The Neural Network model got the greatest values of accuracy and F1-score. This implies that the Neural Network was capable of learning the nonlinear complicated patterns that exist in the dataset, thereby resulting in higher generalisation and improved prediction performance. Conversely, a Decision Tree model exhibited mediocre performance. Although the predictive power was not as good as the Neural Network it was more interpretable in that the decision making process could be followed through its hierarchical structure. However it also showed a tendency to overfit which is a known limitation of tree-based models. The findings of experiments justify significantly the methods chosen in this research. The implementation and comparison of these three models allowed us to verify not only the operational validity of our method but also the best model for the dataset. The comparison research clearly shows that kNN and Decision Tree models give helpful insights and interpretability but the Neural Network surpassed them all in terms of accuracy and F1-score. This

result demonstrates the ability of neural networks to cope with complicated data, and validates their suitability as the core model of our technique.

Table 1: Classification Accuracy

Models Compared	kNN	Neural Network	Tree
kNN	-	0.001 / 0.362	0.001 / 0.749
Neural Network	0.637 / 0.362	-	0.045 / 0.947
Tree	0.251 / 0.749	0.008 / 0.947	-

Table 1 presents the classification accuracy results. The results of three models (kNN, Neural Network, Decision Tree) are compared. The table shows the scores of the accuracy and the values of the statistical significance (p-values). The kNN vs the Neural Network is 0.001 / 0.362. This means that while the Neural Network performed better than kNN, the accuracy of kNN was still rather poor, at 0.362. The figure of 0.001 / 0.749 indicates that Decision Tree achieved a lot better accuracy of 0.749 and greatly improved upon kNN, compared to kNN and Decision Tree. The comparison for Neural Network and Decision Tree offers mixed results. The Neural Network vs the Tree with 0.045 / 0.947. We can see that both models performed quite well but the Decision Tree performed a little better than the Neural Network with accuracy of 0.947. Considering the Neural Network alone, its comparison to kNN again points to its superiority, but not as strong as the Tree. The table generally shows that the Decision Tree had the greatest classification accuracy of the three models, followed by the Neural Network, while the kNN model was lower with limited predictive potential.

Table 2: F1 score

Compare Models	kNN	Neural Network	Tree
kNN	-	0.001 / 0.512	0.001 / 0.718
Neural Network	0.488 / 0.512	-	0.028 / 0.960
Tree	0.281 / 0.718	0.011 / 0.960	-

The F1-score results in Table 2 allow a comparison between the three models, i.e. kNN, Neural Network and Decision Tree. Comparing the kNN with the Neural Network, the value 0.001 / 0.512 indicates that the Neural Network achieved a significantly better F1-score of 0.512 and the kNN was behind. Furthermore, the decision tree was much better than the kNN, with much higher F1-score of 0.718 as shown in the comparison between the kNN and Decision Tree with the number 0.001 / 0.718. The comparison between the neural network and the decision tree shows the number 0.028 / 0.960 with both models performing well but the decision tree performing better and getting the highest F1-score of 0.960. This is also supported by reciprocal comparisons where the neural network obtained a moderate performance against the kNN but behind the decision tree. Overall, the table shows a clear ranking in terms of F1-score, the decision tree showing the best performance, followed by the neural network and the kNN getting the weakest results.

5.1 Fuzzification

In the proposed system, the crisp input values are transformed into fuzzy values through triangular membership functions. Each fuzzy set is represented by three parameters (lo, mi, up), where lo denotes the lower boundary, mi denotes the peak point with maximum membership, and up denotes the upper boundary. The membership value of an input x is calculated using the following formula:

$$u(X) = \begin{cases} \frac{x - lo}{mi - lo} & lo \leq x \leq mi \\ \frac{up - x}{up - mi} & mi \leq x \leq up \\ 0 & \text{otherwise} \end{cases}$$

This function ensures that the degree of membership gradually increases from 0 to 1 as the input approaches the peak value mi, and decreases back to 0 as it moves away from the peak towards the upper boundary up. For a triangular fuzzy set defined by three points (lo, mi, up)

- lo = lower bound where membership starts at 0
- mi = peak point where membership = 1
- up = upper bound where membership falls back to 0

Table 3: Model_1 Range

Fuzzy Set	lo	mi	up
Low	0	0.25	0.5
Medium	0.25	0.5	0.75
High	0.5	0.75	1

Table 3 defines the fuzzification ranges for Model_1, where the input space between 0 and 1 is divided into three overlapping fuzzy sets: Low, Medium, and High. The Low set is represented by the range (0, 0.25, 0.5), meaning values closer to 0 are strongly classified as Low, and membership decreases gradually until 0.5. The Medium set is defined as (0.25, 0.5, 0.75), where the centre value 0.5 has the strongest membership in Medium, and the membership tapers off symmetrically toward 0.25 and 0.75. Finally, the High set is represented by (0.5, 0.75, 1), where values around 0.75 are strongly classified as High, and the membership declines toward 0.5 and 1. This overlapping structure ensures that input values can partially belong to more than one fuzzy set, allowing smooth transitions between categories and making the classification more flexible and realistic.

Table 4: Model_2 Range

Fuzzy Set	lo	mi	up
Low	0	0.33	0.66
Medium	0.25	0.50	0.75
High	0.50	0.75	1

Table 4 shows the fuzzification ranges for Model_2, where the input space from 0 to 1 is divided into three fuzzy sets with slightly wider overlaps compared to Model_1. The Low set is defined as (0, 0.33, 0.66), meaning values starting from 0 are strongly classified as Low, with the highest membership at 0.33, and then gradually decreasing until 0.66. The Medium set is given as (0.25, 0.50, 0.75), where the strongest membership is at 0.5, and it tapers symmetrically toward 0.25 and 0.75. The High set is represented by (0.50, 0.75, 1), where values near 0.75 are classified as High with maximum membership, while membership reduces toward 0.5 and 1. The overlapping between these sets is more significant than in Model_1, meaning that a single input can belong to both Low and Medium, or Medium and High, at the same time. This greater overlap allows Model_2 to reduce sensitivity to small input changes and provides smoother classification boundaries.

Table 5: Model_3 Range

Fuzzy Set	lo	mi	up
Very Low	0	0.10	0.20
Low	0.15	0.30	0.45
Medium	0.40	0.50	0.60
High	0.55	0.70	0.85
Very High	0.80	0.90	1.00

Table 5 presents the fuzzification ranges for Model_3, where the input space from 0 to 1 is divided into five narrower fuzzy sets, giving this model finer granularity compared to Model_1 and Model_2. The Very Low set (0, 0.10, 0.20) captures values close to zero, with maximum membership around 0.10. The Low set (0.15, 0.30, 0.45) shifts slightly upward, focusing on small but not minimal values, with peak membership at 0.30. The Medium set (0.40, 0.50, 0.60) centers around 0.50, representing mid-range inputs with balanced membership. The High set (0.55, 0.70, 0.85) emphasizes larger input values, peaking at 0.70, while the Very High set (0.80, 0.90, 1.00) captures values near the maximum, with peak membership at 0.90. Because the sets are narrower and more numerous, Model_3 allows for more precise classification and smoother transitions between fuzzy categories. This design improves sensitivity, enabling the system to distinguish subtle differences in input values more effectively.




Figure 2: Surface between Model 1 and 2

Figure 2 illustrates the 3D surface plot showing the relationship between Model_1, Model_2 and Accuracy. The x-axis corresponds to Model_1 values ranging from 0 to 1, while the y-axis represents Model_2 values within the same range. The z-axis shows the accuracy, which lies between 0.485 and 0.5. The surface highlights that accuracy does not change linearly with the variation of Model_1 and Model_2; instead, it exhibits a non-uniform pattern. The outer regions of the plot, shown in yellow, indicate higher accuracy values close to 0.5, suggesting that when either Model_1 or Model_2 dominates (values close to the extremes), performance is slightly better. In contrast, the central region, shown in blue, forms a basin where accuracy dips towards 0.485, indicating that when both models contribute equally (mid-range values), the accuracy decreases. The transition regions, represented in cyan and green, demonstrate gradual shifts in accuracy between the high and low zones. Overall, the plot reveals that the combined effect of Model_1 and Model_2 is nonlinear, with the highest accuracy achieved at the boundaries and the lowest at the center, emphasizing the importance of carefully tuning their balance for optimal performance.




Figure 3: Surface between Model 1 and 3

Figure 3 illustrates the 3D surface plot showing the interaction between Model_1, Model_3 and Accuracy. The x-axis represents Model_1 values ranging from 0 to 1, the y-axis corresponds to Model_3 values within the same range, and the z-axis shows accuracy values between 0.49 and 0.52. Compared to the previous figure, the accuracy range here is slightly higher, indicating marginally improved performance when Model_1 and Model_3 interact. The surface shows an irregular, fluctuating pattern with peaks and troughs. Certain regions, particularly in yellow, reach accuracy values above 0.51 and up to 0.52, highlighting combinations of Model_1 and Model_3 where performance improves. Conversely, the blue regions depict dips where accuracy falls closer to 0.49, suggesting less favourable parameter combinations. The presence of sharp peaks indicates that small changes in Model_1 or Model_3 can cause noticeable fluctuations in accuracy, reflecting the sensitivity of the model's behavior to parameter tuning. Overall, this figure reveals that the joint effect of Model_1 and Model_3 is highly dynamic, with some parameter combinations yielding better accuracy, while others lead to performance drops, emphasizing the need for precise tuning to identify optimal regions.

6. Discussion and Conclusion

The findings of this study provide important insights into the development of privacy-preserving machine learning approaches for healthcare data analysis. The comparative evaluation of k-Nearest Neighbors (kNN), Neural Networks, and Decision Tree models demonstrates that model selection plays a significant role in balancing predictive accuracy, interpretability, and robustness. Decision Trees achieved comparatively better accuracy under certain conditions and offered greater interpretability, whereas Neural Networks demonstrated a stronger ability to capture complex and non-linear patterns in healthcare

data. In contrast, kNN served as a useful baseline model but showed relatively limited predictive performance. These observations indicate that the suitability of a model depends on the characteristics of the healthcare application, data distribution, and optimization strategy rather than on the universal superiority of a single algorithm.

The surface plots corresponding to Model_1, Model_2, and Model_3 further demonstrate the dynamic and non-linear relationships among the evaluated parameters. Prediction accuracy varied across different parameter combinations, with some regions showing improvement and others exhibiting reduced performance. These variations emphasize the importance of appropriate parameter tuning, fuzzification, and adaptive optimization for achieving stable and reliable results. In particular, Model_3 demonstrated greater granularity and sensitivity, suggesting that refined fuzzy representations can capture subtle variations in healthcare data and potentially improve predictive performance compared with simpler configurations.

Privacy preservation is another important component of the proposed framework. The integration of differential privacy, federated learning, and secure aggregation provides mechanisms for protecting sensitive healthcare information while enabling collaborative model development. Techniques such as gradient clipping, controlled noise injection, and robust aggregation further strengthen protection against privacy leakage and adversarial manipulation while maintaining predictive utility. These mechanisms also support compliance with healthcare data protection requirements, including HIPAA and GDPR.

The findings further highlight the importance of fairness, transparency, and explainability in healthcare machine learning. Explainable and bias-aware models are particularly important for improving the understanding and acceptance of data-driven decisions by healthcare professionals and patients. In addition, scalability and computational efficiency remain essential considerations for practical deployment, particularly in resource-constrained healthcare environments.

In conclusion, this study demonstrates the potential of integrating machine learning, adaptive optimization, and privacy-preserving mechanisms for secure and effective healthcare data analytics. The results show that model performance depends on algorithm selection, parameter configuration, and optimization strategy, while privacy mechanisms provide essential protection for sensitive health information. Future work should focus on validation using larger and more diverse real-world healthcare datasets, improved explainability and bias mitigation, computational efficiency, and comprehensive evaluation under different privacy and adversarial settings.

Acknowledgment

The authors thank the Department of Computer Applications, Integral University, Lucknow, for providing the necessary support to carry out this research. This work was carried out at Integral University under the reference number IU/ R&D/2026-MCN0004992.

References

1. H. S. Saad, J. F. W. Zaki, and M. M. Abdelsalam, "Employing of machine learning and wearable devices in healthcare system: Tasks and challenges," *Neural Computing and Applications*, vol. 36, pp. 17829–17849, 2024.
2. G. Palozzi and F. Ranalli, "Telemedicine implementation between innovation and sustainability: An operating model for designing patient-centered healthcare," in *Human-Centered Service Design for Healthcare Transformation: Development, Innovation, Change*, Springer, Cham, Springer International Publishing, pp. 375–399, 2023.
3. T. Tadesse, H. Abuye, and G. Tilahun, "Availability and affordability of children essential medicines in health facilities of southern nations, nationalities, and people region, Ethiopia: Key determinants for access," *BMC Public Health*, vol. 21, Art. no. 714, 2021.
4. G. Mattison, O. Canfell, D. Forrester, C. Dobbins, D. Smith, J. Töyräs, and C. Sullivan, "The influence of wearables on health care outcomes in chronic disease: Systematic review," *Journal of Medical Internet Research*, vol. 24, no. 7, Art. no. e36690, 2022.
5. O. Ali, W. Abdelbaki, A. Shrestha, E. Elbasi, M. A. A. Alryalat, and Y. K. Dwivedi, "A systematic literature review of artificial intelligence in the healthcare sector: Benefits, challenges, methodologies, and functionalities," *Journal of Innovation & Knowledge*, vol. 8, no. 1, Art. no. 100333, 2023.
6. S. Aminizadeh et al., "Opportunities and challenges of artificial intelligence and distributed systems to improve the quality of healthcare service," *Artificial Intelligence in Medicine*, vol. 149, Art. no. 102779, 2024.
7. K. Cusi et al., "American Association of Clinical Endocrinology clinical practice guideline for the diagnosis and management of nonalcoholic fatty liver disease in primary care and endocrinology clinical settings: Co-sponsored by the American Association for the Study of Liver Diseases (AASLD)," *Endocrine Practice*, vol. 28, no. 5, pp. 528–562, 2022.

8. M. Masoumian Hosseini, S. T. Masoumian Hosseini, K. Qayumi, S. Hosseinzadeh, and S. S. Sajadi Tabar, "Smartwatches in healthcare medicine: Assistance and monitoring; a scoping review," *BMC Medical Information and Decision Making*, vol. 23, Art. no. 248, 2023.
9. Y. Ye et al., "Management of medical and health big data based on integrated learning-based health care system: A review and comparative analysis," *Computer Methods and Programs in Biomedicine*, vol. 209, Art. no. 106293, 2021.
10. A. Rehman, S. Naz, and I. Razzak, "Leveraging big data analytics in healthcare enhancement: Trends, challenges and opportunities," *Multimedia Systems*, vol. 28, no. 4, pp. 1339–1371, 2022.
11. A. Khanna, P. Selvaraj, D. Gupta, T. H. Sheikh, P. K. Pareek, and V. Shankar, "Internet of things and deep learning enabled healthcare disease diagnosis using biomedical electrocardiogram signals," *Expert Systems*, vol. 40, no. 4, Art. no. e12864, 2023.
12. M. Oloko-Oba, E. Esenogho, and K. Aruleba, "From biosignals to bedside: A review of real-time edge machine learning for wearable health monitoring," *Bioengineering*, vol. 13, no. 5, Art. no. 559, 2026.
13. J. Kumari, E. Kumar, and D. Kumar, "A structured analysis to study the role of machine learning and deep learning in the healthcare sector with big data analytics," *Archives of Computational Methods in Engineering*, vol. 30, no. 6, pp. 3673–3701, 2023.
14. N. Y. Philip, J. J. P. C. Rodrigues, H. Wang, S. J. Fong, and J. Chen, "Internet of Things for in-home health monitoring systems: Current advances, challenges and future directions," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 2, pp. 300–310, 2021.
15. P. Patil, "Healthcare Dataset," Kaggle, 2024. [Online]. Available: Healthcare Dataset — Kaggle. [Accessed: Jun. 5, 2026].