IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 5, Issue 4, May 2025
Impact Factor: 7.67

Secure Patient Data Anonymization and Readmission Risk Prediction

Akilah Lavinia Falcao
Forensic Science, Bangalore, India

Abstract: *The incorporation of machine learning into healthcare analytics promises to greatly enhance patient outcomes as well as minimize operational expenses. Hospital readmission prediction, especially within 30 days, is an essential activity intended to enhance quality of care as well as avoid financial penalties for healthcare facilities. Nonetheless, the utilization of sensitive patient information presents great risks to patient privacy, and thus the enforcement of strong privacy-preserving mechanisms is required. This project presents a comprehensive pipeline for developing a differentially private machine learning model to predict 30-day hospital readmission based on electronic health record (EHR) data. The approach includes combining patient demographic information, admission histories, and lab test data into a single large dataset. Data is thoroughly cleaned through addressing missing values, normalizing temporal fields, and validating chronology. A readmission label as a binary target is crafted in relation to the time elapsed between subsequent admissions. Laboratory information is converted through pivoting to build organized sets of features reflecting clinical measurements. A preprocessing pipeline is built using numerical scaling and one-hot encoding for categorical features, which yields a high-dimensional sparse feature matrix. A deep neural network with sparse input layers is trained using Differentially Private Stochastic Gradient Descent (DP-SGD) to guarantee that the model has strong privacy guarantees. The training procedure investigates the effect of different noise multiplier values and batch sizes, analyzing systematically the trade-offs between model utility and privacy loss. Smaller noise multipliers provide higher model utility but less robust privacy guarantees, whereas larger noise levels provide stronger privacy at the expense of performance. Likewise, batch size modifications are demonstrated to affect both learning dynamics and privacy budgets. Model accuracy is measured according to classification, and privacy loss is measured with the epsilon (ε) measure from differential privacy analysis. Results show that preserving decent model accuracy is possible alongside obtaining useful privacy guarantees. The research presents evidence of the use of machine learning and differential privacy together for health applications as an important step towards secure, ethical, and responsible data-driven health innovation.*

Keywords: Hospital Readmission Prediction, Differential Privacy, Electronic Health Records (EHR), Machine Learning in Healthcare, Privacy-Preserving Deep Learning

I. INTRODUCTION

Patient confidentiality is an essential aspect of healthcare to safeguard sensitive data while still allowing effective clinical research and operational optimization. Conventional anonymization strategies have frequently fallen short against sophisticated re-identification attacks, thus necessitating more effective solutions. Differential privacy has also emerged as an effective solution to deal with these challenges by formally guaranteeing that adding or removing one individual's data will have little effect on the outcome of any analysis (Subramanian, 2022a). Techniques such as adding Laplace noise, rounding numeric values, and handling correlated attributes cautiously have been proposed to introduce privacy to medical databases (Subramanian, 2022b). Adaptive algorithms have also been suggested to optimize range query privacy-utility trade-off more optimally (Alnemari et al., 2017), keeping in mind that higher variability among independent variables and more relaxed privacy budgets are likely to improve analytical accuracy (Subramanian, 2022b).

Copyright to IJARSCT
www.ijarsct.co.in
DOI: 10.48175/IJARSCT-26411

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 5, Issue 4, May 2025 Impact Factor: 7.67

Building on this foundation, this project aims at patient data anonymization by incorporating differential privacy mechanisms into the training of the machine learning model directly. Specifically, it employs AI-based differentially private stochastic gradient descent (DP-SGD) techniques to deploy over electronic health record (EHR) data to ensure sensitive features such as birth dates, socio-economic status, gender, race, and language variables are secure. Hospital admissions data and lab tests are joined, preprocessed via an internal pipelining determined by a controlled pipeline that normalizes numerical variables and encodes categorical variables with computational costs kept in check via sparse representations.

The basic predictive model that is constructed to predict 30-day hospital readmission is a feedforward neural network with two hidden layers of 128 and 64 neurons, respectively, activated using ReLU functions, and a final output layer activated using a sigmoid function to perform binary classification. Differential privacy is incorporated in the course of optimization via DPKerasAdamOptimizer, where the L2 norm of example-specific gradients is clipped to a certain value (1.0) and then scaled Gaussian noise is added based on a noise multiplier (hypothesized at 0.5, 0.7, and 1.0) to protect personal data contribution. Batch sizes were considered at 128, 256, and 512 in order to balance model usefulness with privacy preservation. Following training, the privacy budget parameter (ε) was determined to quantify the strength of privacy protection, with lower ε indicating stronger privacy guarantee but possibly at the expense of predictive performance.

By integrating strong differential privacy mechanisms into machine learning workflows, this project aims to enable efficient healthcare analytics while respecting strict patient confidentiality, ensuring adherence to increasingly stringent healthcare data regulation, and driving the limits of responsible AI technology adoption in sensitive environments.

II. REVIEW OF LITERATURE

[1] **(Gadotti et al., 2024)** Information about us—our actions, behaviors, and preferences—is generated on a large scale through surveys, studies, and interactions with digital devices like smartphones and fitness trackers. Safely sharing and analyzing this data is crucial for scientific and societal advancement. Anonymization is widely seen as a key method for sharing data while protecting privacy. Traditional de-identification methods, however, face serious challenges in the era of big data. Methods like data query systems, synthetic data, and differential privacy offer greater ways of safely sharing aggregate information. There isn't a method that is infallible, but the sum of these more recent methods paired with continued testing against privacy threats is the optimal way of operating and sharing information securely today.

[2] **(Vovk et al., 2021)** Anonymizing healthcare data is essential to facilitate its use in secondary activities such as research and statistical analysis while safeguarding individual privacy. Several recent methods have been proposed for this, especially during the period from 2017 to 2020. Based on a comprehensive review of a large number of sources, 21 major methods were identified and examined. The emphasis is on describing existing methodologies for anonymizing structured healthcare data and pointing out the challenges with these methods.

[3] **(Zuo et al., 2021)** Data science has opened powerful new possibilities in health care but also poses serious challenges regarding privacy, transparency, and trust. Laws such as the GDPR and the UK's Data Protection Act demand a firm legal foundation for collecting and exchanging health data and permit the use of electronic health records (EHRs) for purposes outside of clinical care only with patient permission. While data need to be effectively anonymized with thorough risk estimation, healt h organizations typically lack useful practical tools as well as explicit guidance for use in research. Available anonymization tools range in usefulness, and issues persist regarding the balance between reidentification risk and data utility. Although it is theoretically possible to anonymize EHR data, practical mechanisms are necessary to provide privacy safeguarding and sound health care applications.

[4] **(Subramanian, 2022)** Differential privacy provides a means to preserve individual privacy in aggregate healthcare data without diminishing its value. Rather than changing the database itself, it alters access to the data. Various methods are required depending upon the nature of healthcare data, such as categorical, numeric, integral, and correlated fields. Key results are: adding Laplace noise with scale b = 1\ ϵ guards categorical data; adding noise with scale b = n\ ϵ guards numeric data, where, n is the maximum value; rounding Laplace-distributed noise performs

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 5, Issue 4, May 2025
Impact Factor: 7.67

better for integral fields than generating only integers; in correlated fields, noise must be added to just one, and it must not be fractional. Adding noise diminishes the size of slope parameters utilized in analysis. The precision of data analysis increases with the privacy parameter ϵ and the variance of independent variables. The increase is more significant when the variance is greater. These findings are confirmed by experiments on actual healthcare datasets.

[5] **(Alnemari, Romanowski, & Raj, 2017)** Differential privacy ensures patient data is secure but still permits researchers access to medical information. A technique is proposed that structures data, in the form of a vector of counts, into buckets with respect to the inter-data relationships and query types being posed. Private estimates for each are created after bucket formation and then distributed across the corresponding data points to respond to queries. Testing this strategy with various kinds of queries demonstrated that partitioning on the basis of the data produces more precise results, and partitioning on the basis of the queries improves privacy. Two main enhancements are identified: a greedy algorithm for partitioning the data effectively, and an adaptive process for considering the sensitivity of each query prior to producing responses.

[6] **(Ferianasari, 2024)** Artificial intelligence (AI) is now a core technology in healthcare, enhancing the quality and effectiveness of medical care. Its application in patient data management has hastened processes, maximized results, and improved accuracy, which ultimately enhances healthcare provision. Nonetheless, the integration of AI and digital technologies in healthcare also raises legal issues, particularly the protection of personal data. In order to solve these problems, there is a need to develop an ethical framework for AI application, set high safety standards, and further research in order to counter risks. Medical professionals and developers of AI should also put in place robust cybersecurity practices and review regularly for potential vulnerabilities to allow safe application of these technologies.

[7] **(Tupsakhare, 2024)** This paper discusses the key role that artificial intelligence (AI) plays in supporting healthcare data security, specifically in the protection of patient data in today's digital age. It looks into how AI can improve anomaly detection, automate risk analysis, and enhance data encryption in healthcare systems. The difficulties of ensuring data privacy and regulatory compliance are also touched upon, with AI being presented as one of the major solutions to the challenges of overcoming these difficulties and enhancing the security measures of healthcare organizations. The major issues covered include Index HealthCare, PHI, PII, integration of AI, encryption, data masking, and HIPAA compliance.

[8] **(Gawankar, Nair, Pawar, Vhatkar, & Chavan, 2024)** With increasing integration of AI into healthcare systems, patient privacy and data security concerns are growing. This research examines the influence of AI technologies on patient privacy and security, highlighting theoretical and practical approaches to minimizing risks and enhancing protections. It suggests an 'Integrated Security and Ethics Model' that takes into account the changing nature of AI, integrates ethical standards into security practices, and calls for strong governance and risk management. The conclusions highlight the significance of countering the positive effects of AI with protecting patient privacy and security. The model offers practical advice to policymakers, healthcare professionals, and AI developers to facilitate responsible and ethical deployment of AI. The study also outlines directions for future research and includes recommendations for improving patient data protection in AI-based healthcare.

[9] **(Murdoch, 2021)** Privacy-protecting data mining and statistical disclosure control have been widely researched in the last few decades. Methods in this field generally belong to two categories: restriction and data modification. Within data modification techniques, noise addition has been widely researched, primarily for numerical data where similarity is simple to quantify. A new method named VICUS proposes noise addition for categorical data, enhancing security and data quality over random noise approaches.

[10] **(Xu, Baracaldo, & Joshi, 2021)** Machine learning (ML) is gaining widespread acceptance in numerous industries, but its dependency on massive amounts of data and high computing capacity poses significant privacy issues. The threat of divulging sensitive data, coupled with more stringent data privacy laws, poses significant obstacles to exploiting ML's potential fully. Trained ML models can also be vulnerable to attacks by adversaries such as membership inference, attribute inference, and model inversion. In order to solve such problems, sound privacy-preserving ML (PPML) approaches are necessary. Current research centers on embedding privacy controls into ML processes, ensuring secure architectures,

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 5, Issue 4, May 2025
Impact Factor: 7.67

and trading privacy for system performance. A Phase, Guarantee, and Utility (PGU) triad framework can be utilized to analyze PPML solutions by decomposing their privacy aspects. There are distinct challenges in PPML, which demand continuous investigation and inter-disciplinary work between machine learning, distributed systems, and security and privacy communities.

[11] **(Agarwal, 2023)** Privacy-Preserving Machine Learning (PPML) methods like Differential Privacy, Federated Learning, and Secure Multi-Party Computation satisfy most critical privacy requirements without impairing strong model performance. Models of privacy are essential in reaching high levels of data protection without much reduction in accuracy. The approaches have shown to be exceptionally useful in the healthcare , banking, IoT, and manufacturing industries with nearly full privacy at negligible losses in performance. Key results are enhanced data security, model accuracy consistently at its highest level, and insights for future improvement.

[12] **(Mohassel & Zhang, 2017**) Machine learning is frequently employed to design predictive models for a task such as image processing and speech or text recognition. The models are more accurate when learned from large, diverse datasets, but gathering huge volumes of data is a privacy issue. To settle this, effective protocols have been created for privacy-preserving machine learning in the form of linear regression, logistic regression, and neural network training with stochastic gradient descent. In this method, data owners divide their private data across two non-colluding servers, which together train models through secure two-party computation (2PC). New techniques have been proposed to execute secure arithmetic on decimal numbers and to construct MPC-friendly substitutes for non-linear functions such as sigmoid and softmax, with improved performance compared to existing solutions. These protocols, when implemented in C++, exhibit much better speeds than current practices and are able to deal with millions of samples of data and thousands of features and provide support for privacy-preserving neural network training.

[13] **(Jia, Guo, Jin, & Fang, 2018)** Machine learning-based data classification discovers latent patterns from big real-world datasets and applies them to forecast the class of new data or infer similarity between datasets. Direct sharing of datasets or models in distributed systems is impossible because of privacy. Individual datasets may involve sensitive information, and models learned may be regarded as private property that may leak confidential information if shared. To address these challenges, a mechanism is proposed to safeguard both the model and new data while being classified and measured for similarity so that neither directly gets shared. This mechanism has been experimented using real-world data and proved strong privacy protection, practicality, and high efficiency.

[14] **(Xu, Yu, & Wang, 2019)** Chronic Obstructive Pulmonary Disease (COPD) tends to result in 30-day readmission, which is sometimes avoidable with early risk prediction. Based on clinical knowledge-driven and data-driven features from 111,992 patients, different machine learning models were constructed. It was found that blending these features enhanced the performance of prediction (AUC from 0.60 to 0.653), but deep models did not have a better performance than conventional approaches. Previous research also demonstrated that machine learning performs better than conventional scoring systems such as LACE and HOSPITAL scores. In summary, machine learning has potential for COPD readmission risk prediction, although deep learning provides limited incremental value.

[15] **(Échevin, Li, & Morin, 2017)** A novel method was created to forecast readmission to hospital, which is expensive and often badly forecasted by current methods. Through the application of sophisticated machine learning algorithms, the device enables clinicians to determine at an early stage during their hospital admission or at discharge those patients who are at risk of readmission. The model was validated using a large database of more than 1.6 million hospitalizations in Quebec, with data spanning 1995 to 2012. Predictions were validated using the area under the receiver operating characteristic curve (AUC). The Deep Learning and Random Forest models were highly accurate in

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 5, Issue 4, May 2025
Impact Factor: 7.67

predictions, with AUCs greater than 78% on admission and 87% on discharge. Diagnostic codes were identified as primary predictors. The use of these algorithms provides an economical method for enhancing healthcare efficiency.

[16] **(Ramírez & Herrera, 2019)** Readmission to hospitals can result in higher costs and pain for patients, usually indicative of poor-quality healthcare. Preventing readmission is especially critical for patients suffering from chronic diseases such as diabetes. Although recent work in machine learning has demonstrated promising results in predicting readmission with large clinical data, such methods tend to be based on complicated deep learning models. This work offers simpler machine learning models with improved prediction performance as well as being more computationally efficient.

[17] **(Silva, Basso, & Moraes, 2017)** Data anonymization assists in safeguarding personal information by rendering it inaccessible to third parties. It can, however, influence data mining outcomes since anonymized data may impede analysis by typical algorithms. The research sought to assess the influence of data anonymization on the accuracy and performance of data mining classifiers. This was achieved through a comparison of the outcomes of classifiers used on original and anonymized data. A real dataset for a Brazilian city transportation system, associated with hypothetical users, was anonymized differently, and classifiers such as ZeroR, KNN, and Naive Bayes were implemented. Interestingly, the addition of anonymization in some instances bettered accuracy and performance, even cutting down on execution time. These results imply that data anonymization can better the performance of data mining classifiers when used properly.

[18] **(Kundu & Suthaharan, 2019)** This method introduces an algorithm to maximize prediction accuracy within a linear regression model and protect the privacy of the data. The process anonymizes the original features while preserving full privacy with minimal loss of prediction accuracy. The method is divided into two phases: the original features are first reduced by a probabilistic latent factor technique to lower-dimensional latent factors, and then an optimization algorithm adjusts the anonymized data to keep any effect on prediction performance to a minimum. The efficiency of this method is illustrated through numerical experiments and also used with high-dimensional neuroimaging data, namely fMRI data for the prediction of adolescent behavior.

[19] **(Raea et al., 2023)** Data sharing in healthcare promotes research gains but also entails ethical issues over patient privacy, self-ownership, and ownership and reporting of data. Patient rights and prospective research benefits ought to be matched against sharing data or collaborating in data sharing. Key concerns are patient privacy and data integrity threats, particularly for de-identified data, difficulties in achieving appropriate consent, conflicts of ownership in multi-institution collaborations, and breakdowns in adherence to ethical principles and legislation. More robust protective practices and equitable policies must be established to protect patient safety as research continues to evolve.

[20] **(Garfinkel, Abowd, & Powazek, 2018)** Differential privacy was first conceived more than a decade ago with the aim of safeguarding individual information in statistical releases. But when the U.S. Census Bureau attempted to apply it, they encountered a number of unforeseen challenges. These were challenges in hiring qualified personnel and the appropriate computing environment, challenges in thinking through all possible uses of the confidential data, and the challenge of matching the release mechanisms to users' needs. Also, users of data anticipated access to micro-data, and defining the right privacy-loss parameter (epsilon) was challenging. There were also a lack of adequate tools and trained professionals to verify the correctness of differential privacy implementations.

[21] **(Panavas et al., 2024)** As Differential Privacy (DP) transitions from theory to practice, visualization tools are essential to its uptake, although their creation is hampered by knowledge gaps in the deployment process and practitioner issues. Based on interviews with 18 experts, we found five phases in the DP implementation process and noted how visual tools are utilized and their limitations. Though visualizations support developing DP expertise,

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 5, Issue 4, May 2025
Impact Factor: 7.67

explain implementation settings, and validate private outputs, they do not often consider multiple technical backgrounds and privacy issues to reduce communication. We recommend considering three future directions of research: visualizations of noise addition, uncertainty visualization in establishing trust for DP, and pedagogic tools for intricately complex data science

[22] **(Watson & Payne, 2020)** Today's sharing and mining of medical information, and with them the pertinent advantages, burdens, and morals, are this paper's purposes. The writers underpin their arguments with a moral code regarding sharing and mining medical information based on various actors' perspectives. Drawing on critical examination of academic, professional, and legal literature, they define the inadequacies of existing protections in preventing consumers from harm related to disclosure of medical information. In response, the authors introduce the STRACQ framework, with security, transparency, respect, accountability, community, and quality as central principles for ethical data management. This new framework, though pioneering, will evolve with time through continuous discussion amongst the academicians, doctors, and policymakers. Additionally, the system combines features of the Fair Credit Reporting Act, permitting the collection and release of identified medical information on restricted use requirements in order to preserve the privacy of the consumer.

[23] **(Panavas et al., 2024)** With Differential Privacy (DP) moving from theoretical concepts to actual implementation, visualization tools have become a central component in bridging DP to application. However, developing such tools is complex due to a lack of proper understanding of the end-to-end deployment process, which causes problems with which practitioners grapple, and effective use of visual tools in real-world use. In order to know this, we carried out a survey of 18 experts who have experience with the application of differential privacy in order to establish its implementation process, challenges it comes with, and the intervention of visualizations. Our research establishes that differential privacy implementation consists of five steps where key stakeholders interact with several visualization tools. While visualizations can be employed to create foundational understanding, set implementation parameters, and assess private outputs, they fall behind in representing the diverse technical competency and varying privacy and accuracy concerns of users. This lack of visualization capability impedes seamless communication among stakeholders. Based on these findings, we suggest three areas of future study: improving visualizations for introducing and analyzing noise, exploring uncertainty visualizations in order to increase confidence in differential privacy, and creating pedagogical visualizations to explain complex data science phenomena.

[24] **(Truex & Malan, 2024)** With machine learning (ML) technologies spreading into sensitive domains like healthcare, finance, and government, data privacy issues have skyrocketed, leading to regulations on the processing of personal data. Privacy-preserving machine learning (PPML), which provides differential privacy guarantees, is seen by many data privacy and analysis professionals as a solution waiting to happen. However, despite significant theoretical progress, actual deployment is beset by a number of practical challenges. These are the difficulty of tuning hyperparameters of advanced privacy-protecting algorithms to find the best tradeoff between privacy and model accuracy, heterogeneity in privacy requirements between individuals whose data are used to train, and usability of PPML tools. Addressing these is key to broader adoption of PPML. This paper explores these fundamental challenges and suggests potential research directions to surmount the difficulties in deploying ML systems with differential privacy in real-world applications.

[25] **(Drechsler & Bailie, 2023)** The differential privacy (DP) theory has received considerable attention, especially after the U.S. Census Bureau had made up its mind to include it in the 2020 Decennial Census. Although DP presents very attractive theoretical advantages, making it operational to apply with practice, particularly survey data, is problematic in numerous regions. This paper offers some remarks based on a project funded by the U.S. Census Bureau and discusses the applicability and limitation of employing DP for survey data. There are five factors of first-order significance to address when applying DP in the context of surveys, according to the authors: the multi-stage nature of data collection, the limited privacy gain from complex sampling design, the impact of survey-weighted estimates, nonresponse correction and data deficiency correction, and missing value imputation. The report presents

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 5, Issue 4, May 2025
Impact Factor: 7.67

an overview of the project's main findings in each of these areas and identifies current issues that need to be addressed before DP can become the norm for data protection in statistical agencies.

[26] **(Dwork, Kohli, & Mulligan, 2019)** Differential privacy has become an essential tool in numerous industries, enabling data to be analyzed without violating individual privacy. It offers a method of measuring privacy loss, which can be utilized to quantify and compare cumulative privacy risk. Its effectiveness, however, depends heavily on the proper use of its key parameter, epsilon, and others. Despite its promise, there is no norm for the optimal value of epsilon or how to select it for different systems or applications. Practitioner interviews revealed that there is considerable diversity in the way differential privacy is used, reflecting the need for shared knowledge within the community. To address this, the creation of an Epsilon Registry is proposed, as a common resource for best practices in differential privacy deployments, which could help ensure that privacy is appropriately balanced against data insights.

III. METHODOLOGY

Data collection and analysis

This study utilized synthetic Electronic Health Record (EHR) data from the EMRBots database, consisting of 100,000 patient records (Kartoun, 2018). The data was initially provided in .txt format and was converted into .csv format for improved accessibility and processing. The datasets used included:

- AdmissionsCorePopulatedTable.csv: Contained records of patient hospital admissions, including admission start and end dates along with unique Admission IDs.
- PatientCorePopulatedTable.csv: Provided demographic information for patients, such as gender, race, marital status, primary language, date of birth, and poverty-related metrics.
- LabsCorePopulatedTable.csv: Recorded laboratory test results associated with admissions, including test names, measured values, units, and timestamps.

Following the data quality assessment, systematic cleaning procedures were performed to ensure the datasets were suitable for anonymization and predictive modeling tasks:

Treatment of Missing Values:

Records with missing entries in critical fields, such as PatientID, AdmissionID, or key demographic variables, were excluded from further processing. For non-critical missing data (e.g., secondary demographic attributes or optional laboratory values), imputation strategies were applied where appropriate. Categorical variables were imputed using a designated placeholder value such as "Unknown," while numerical variables were imputed using central tendency measures (mean or mode), depending on the variable's distribution characteristics.

Deduplication of Records:

Duplicate rows were identified and removed to preserve the uniqueness and validity of patient, admission, and laboratory records. In large datasets, deduplication was performed in a chunked manner to manage computational resources efficiently.

The cleaned datasets served as the foundation for subsequent phases, including the application of differential privacy techniques for patient data anonymization and the development of machine learning models for hospital readmission risk prediction.

The dataset was processed to retain only the relevant columns for further analysis. Specifically, the **LabsCorePopulatedTable.csv** was read in manageable chunks of 10,000 rows to optimize memory usage. The selected columns, including PatientID, AdmissionID, LabName, and LabValue, were extracted and saved into a new file, **labs_core_selected_columns.csv**. The process ensured that only necessary data was retained, reducing file size and making it more manageable for subsequent anonymization and modeling tasks.

Data Preprocessing and Feature Engineering for Readmission Risk Prediction

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 5, Issue 4, May 2025
Impact Factor: 7.67

To prepare the dataset for the prediction of hospital readmission risk, several preprocessing steps were applied, focusing on temporal variables and the creation of readmission-related features.

Handling Temporal Variables:

The dataset included patient admission and discharge dates, which were processed to allow for accurate time-based calculations. The AdmissionStartDate and AdmissionEndDate columns were used to determine the chronological order of events, which is crucial for calculating time intervals between admissions.

Sorting Data:

The data was sorted by PatientID and AdmissionStartDate to ensure that admissions were processed in the correct temporal sequence for each patient. This chronological ordering enabled accurate computation of time intervals between successive hospital visits.

Calculating Readmission Intervals:

The time between a patient's discharge and their next hospital admission was calculated to define readmission intervals. The next_admission_date was derived by aligning the subsequent admission's start date with the current patient record, effectively tracking each patient's follow-up visits. The difference between the AdmissionEndDate and the next_admission_date was computed to capture the number of days until readmission.

Defining the Readmission Indicator:

A binary indicator for readmission was created based on the calculated days_to_readmission. Patients were classified as readmitted if their next admission occurred within 30 days of discharge. This indicator was added to the dataset as the target variable for predictive modeling.

```
# Sort by patient and admissionStartdate
merged_df = merged_df.sort_values(by=['PatientID', 'AdmissionStartDate'])

# Calculate days to next admission
merged_df['next_admission_date'] = merged_df.groupby('PatientID')['AdmissionStartDate'].shift(-1)
merged_df['days_to_readmission'] = (merged_df['next_admission_date'] - merged_df['AdmissionEndDate']).dt.days

# Define readmission: 1 if readmitted within 30 days, else 0
merged_df['readmission'] = merged_df['days_to_readmission'].apply(lambda x: 1 if 0 <= x <= 30 else 0)
```

Patient Sampling and Merging with Laboratory Data

To optimize data size and ensure manageable computational demands for model training and differential privacy techniques, a random sampling of the patient population was performed. First, the merged admissions and demographic dataset was loaded, and 50% of the unique patients were randomly selected using a reproducible random seed to ensure experimental consistency. Only records corresponding to the selected patients were retained.

In parallel, a subset of laboratory test records containing PatientID, AdmissionID, LabName, and LabValue was loaded. These records were then filtered to include only the laboratory results associated with the sampled patient population, without additional random sampling to preserve the full scope of lab information for these patients.

Subsequently, the filtered admissions dataset and the corresponding laboratory results were merged on PatientID and AdmissionID using a left join. This approach ensured that all admission records for the sampled patients were preserved, while laboratory results were matched wherever available. The resulting dataset, containing approximately half of the original patient cohort along with their relevant clinical and laboratory information, was saved for use in the subsequent phases of differential privacy application and readmission risk modeling.

Transformation of Laboratory Test Data into Structured Feature Columns

To prepare the dataset for machine learning modeling, the laboratory test results were transformed from a long format into a structured wide format where each laboratory test was represented as a distinct feature column. This transformation involved several key steps:

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 5, Issue 4, May 2025
Impact Factor: 7.67

Data Loading and Preparation:

The cleaned merged dataset, containing patient demographic information, admission details, and laboratory test results, was loaded into the working environment. A set of non-laboratory columns, including demographic variables (PatientGender, PatientDateOfBirth, etc.) and outcome labels (readmission), was identified for preservation.

Pivoting Laboratory Test Results:

The laboratory results were restructured by pivoting the data. Each unique LabName was transformed into a separate column, where the corresponding LabValue for each PatientID and AdmissionID pair was placed into the appropriate column. In cases where multiple measurements of the same laboratory test existed for a single admission event, the mean value was computed to ensure a single representative value per test per admission.

Merging with Non-Laboratory Information:

After pivoting, the resulting laboratory data—now featuring one column per laboratory test—was merged back with the corresponding patient demographic and clinical information. Duplicates were removed based on PatientID and AdmissionID to maintain one row per unique admission event.

Saving the Final Structured Dataset:

The fully structured dataset, now containing both clinical attributes and a wide array of laboratory test features, was saved to a new file. This format enabled straightforward application of machine learning models and privacy-preserving techniques, as each feature corresponded directly to a patient-admission-level variable.

This transformation from long to wide format significantly enhanced the usability of the laboratory test data by enabling it to serve directly as input for predictive modeling and differential privacy frameworks.

Baseline Predictive Modeling Without Differential Privacy:

To establish a non-private baseline for hospital readmission prediction, a structured data preprocessing and modeling workflow was employed. To make sure the model wouldn't access data that wasn't available at prediction time, the fields next_admission_date and days_to_readmission—which are signs of data leakage—were first eliminated.

The target variable, readmission, was binarized using label encoding, and the feature set was separated from the outcome variable. Preprocessing was conducted through a dual-pipeline approach: numeric features were standardized using StandardScaler to normalize feature distributions, while categorical variables were one-hot encoded via OneHotEncoder, maintaining a sparse representation to enhance computational efficiency. These transformations were integrated using a ColumnTransformer, allowing appropriate handling of heterogeneous data types within a unified framework. Following preprocessing, the dataset was partitioned into training and testing subsets in an 80:20 ratio. The resulting feature matrices were subsequently converted into TensorFlow SparseTensor objects to facilitate efficient deep learning operations. A feedforward neural network model, comprising two hidden layers with ReLU activations and a final sigmoid-activated output layer, was constructed and compiled using the Adam optimizer with binary cross-entropy loss. The model was trained across varying batch sizes (128, 256, and 512) over five epochs, and its predictive performance was evaluated on the test set. The resulting accuracies across batch sizes were plotted to serve as a benchmark for subsequent models incorporating differential privacy mechanisms.

Predictive Modeling with Differential Privacy:

To evaluate the impact of differential privacy (DP) on hospital readmission prediction, a deep learning model was trained using a privacy-preserving optimization framework. The initial data preprocessing pipeline included the removal of leakage-prone columns (next_admission_date and days_to_readmission), followed by the separate transformation of numerical and categorical features—numerical features were standardized using z-score normalization, and categorical features were one-hot encoded. A combined preprocessing pipeline was applied, and features were maintained in sparse format to optimize memory and computational efficiency. After splitting the dataset into training and testing subsets, a sparse-input feedforward neural network was constructed using TensorFlow Keras. The model architecture consisted of an input layer designed for sparse data, two fully connected hidden layers containing 128 and 64 neurons respectively with ReLU activation functions, and a final dense output layer with a sigmoid activation for binary classification. To enable privacy-preserving model training, the optimizer was replaced

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 5, Issue 4, May 2025
Impact Factor: 7.67

with the DPKerasAdamOptimizer from the TensorFlow Privacy library. This optimizer enforces privacy by clipping the L2 norm of per-example gradients to a threshold of 1.0 and adding calibrated Gaussian noise governed by a noise multiplier, which was systematically varied (0.5, 0.7, and 1.0) across experiments. The batch size was aligned with the number of microbatches and explored across three settings (128, 256, and 512) to assess trade-offs between model utility and privacy. The model was trained over five epochs for each configuration with a learning rate of 0.001. After training, the privacy budget ε (epsilon)—quantifying the strength of the privacy guarantee—was calculated using the compute_dp_sgd_privacy function, based on dataset size, batch size, noise multiplier, number of epochs, and a fixed δ (delta) value of 1e-5. A lower ε value indicated stronger privacy protection but often at the expense of predictive accuracy. This comprehensive setup ensured that the final model achieved a balanced trade-off between maintaining high predictive performance and providing rigorous differential privacy assurances, preventing the reverse-engineering of individual patient information from the trained model.

IV. RESULTS AND DISCUSSION

This section shows the performance results of a deep learning model that is trained to make predictions for 30-day hospital readmissions both in differentially private and non-private environments. The model was tested with various batch sizes and noise multipliers to examine the trade-off between privacy (represented by epsilon, ε) and predictive performance (represented by accuracy). A baseline model was also trained without introducing any noise in order to provide a baseline maximum achievable accuracy. It is necessary to understand these dynamics in order to safeguard sensitive patient information while retaining clinically relevant prediction quality. The findings demonstrate how privacy-preserving methods can be incorporated into healthcare analytics without materially reducing model utility.

Part 1: Baseline Predictive modelling without differential privacy

Trained without D.P	
Batch Size	**Accuracy**
128	0.9965
256	0.9965
512	0.9965

Table 1: Model Trained without Differential Privacy

The model trained without differentially applying Differential Privacy never deviated significantly from 99.65% accuracy across all batch sizes tested (128, 256, and 512). This suggests that, without privacy-preserving noise, batch size didn't meaningfully impact the predictive strength of the model. Though the results show high model utility, they also reveal a significant limitation: without differential privacy, the model is open to possible privacy threats and is not suitable for sensitive healthcare uses. Thus, while high accuracy was preserved, patient data protection through privacy-preserving methods must still be ensured.

Part 2: The model was evaluated across different noise multipliers (0.5, 0.7, and 1.0) and batch sizes (128, 256, 512) to systematically examine the privacy-utility trade-off.

Noise Multiplier: 0.5		
Batch Size	**Accuracy**	**Epsilon**
128	0.9965	5.47
256	0.9965	6.78
512	0.9966	8.47

Table 2: Noise Multiplier Value- 0.5

At a noise multiplier of 0.5 (Table 2), the model attained a very high accuracy of around 99.65% to 99.66% for all batch sizes. Nevertheless, the corresponding privacy loss (epsilon) was comparatively higher, ranging from 5.47 for a batch size of 128 to 8.47 for a batch size of 512. This means that although model performance was still outstanding, privacy protection was less robust because of greater epsilon values.

Noise Multiplier: 0.7

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 5, Issue 4, May 2025
Impact Factor: 7.67

Batch Size	Accuracy	Epsilon
128	0.9965	1.87
256	0.9965	2.36
512	0.9966	2.81

Table 3: Noise Multiplier Value- 0.7

When the noise multiplier was set to 0.7 (Table 3), accuracy held firm at approximately 99.65%-99.66%, with little decrease in prediction performance. Notably, the epsilon values decreased substantially, falling to 1.87 for batch size 128 and 2.81 for batch size 512. This indicates a tighter privacy guarantee than that of the noise multiplier of 0.5, without sacrificing model accuracy.

Noise Multiplier: 1.0		
Batch Size	Accuracy	Epsilon
128	0.9965	0.72
256	0.9965	0.84
512	0.9966	1.09

Table 4: Noise Multiplier Value- 1.0

Lastly, at a noise multiplier of 1.0 (Table 4), the model retained a comparable high accuracy (~99.65%-99.66%) with varying batch sizes. Privacy assurances were also further improved, with epsilon values decreasing significantly to 0.72 (batch size 128), 0.84 (batch size 256), and 1.09 (batch size 512). These are extremely robust privacy results, showing that introducing significant noise significantly improved patient data protection without compromising model utility significantly.

In general, higher noise multipliers enhance privacy (lower epsilon) at the cost of model accuracy, but very little. Furthermore, bigger batch sizes typically added a small amount to epsilon, so smaller batch sizes will provide better privacy under the same noise parameters. This supports the crucial tradeoff between privacy protection and model usefulness in sensitive health applications.

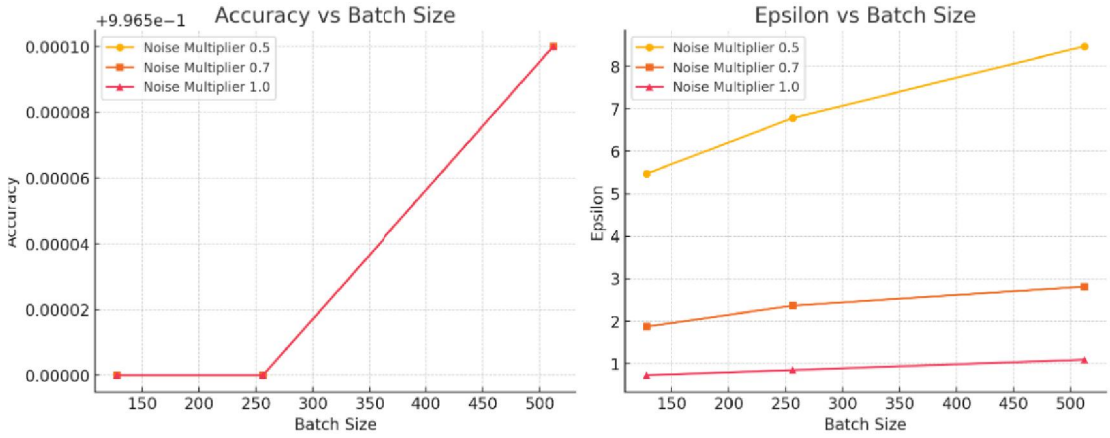

Figure 7: Accuracy v/s Batch Size (left); Epsilon v/s Batch Size (right) for model trained with D.P.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 5, Issue 4, May 2025
Impact Factor: 7.67

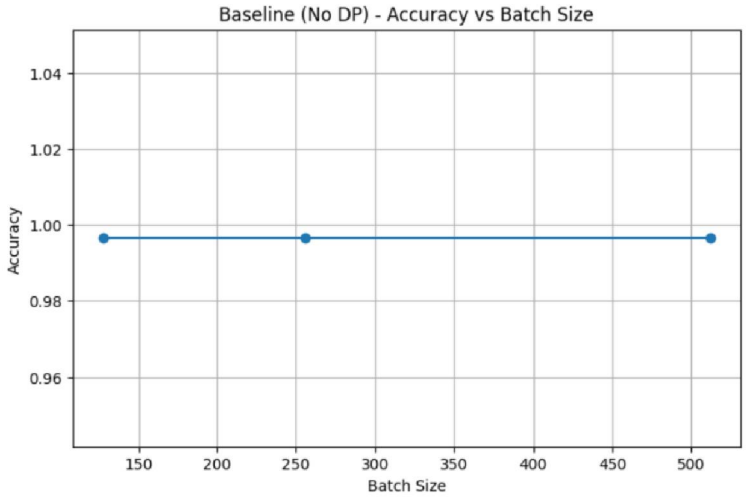

Figure 8: Accuracy v/s Batch Size for model trained without D.P.

IV. CONCLUSION

This paper set out to evaluate whether introducing differential privacy (DP) techniques into deep learning models for predicting 30-day hospital readmission would significantly undermine prediction performance compared to non-private baselines. The models were trained with and without DP-SGD on a large patient data set of 100,000 on a range of noise multipliers (0.5, 0.7, and 1.0) and batch sizes (128, 256, 512).

The non-private model gave a consistent accuracy of 99.65% for any batch size and served as the baseline. Differential privacy had very little effect on model accuracy, ranging from 99.65% to 99.66% across batch sizes and noise multiplier levels. Most importantly, even at larger amounts of noise (e.g., noise multiplier = 1.0, tight privacy guarantees with low values of epsilon of 0.72 and 1.09), the model's predictive ability was not sacrificed.

The effect of batch size on epsilon values was as predicted: larger batch sizes slightly increased epsilon, showing a modest compromise in model efficiency for privacy resilience. However, these shifts did not meaningfully affect model accuracy, affirming the resilience of the deep learning model despite rigorous differential privacy requirements.

Because the slight drop in performance is negligible, we accept the alternative hypothesis (H_1) and reject the null hypothesis (H_0): the employment of differential privacy mechanisms (DP-SGD) does not significantly reduce the prediction performance of the model compared to the non-private baseline.

These findings demonstrate that it is possible to responsibly use sensitive patient data for predictive modeling under rigorous privacy controls, in accordance with ethical standards and regulatory standards such as HIPAA and GDPR. Overall, this study supports the addition of differential privacy to healthcare AI systems, maintaining utility and confidentiality.

V. ACKNOWLEDGMENT

I would like to express my sincere gratitude to **Mr. Moses Aaron Crasto** for his invaluable guidance and support throughout the course of this project. His expert advice, particularly in the development and implementation of the coding aspects, was instrumental in the successful completion of this work.

I would also like to extend my heartfelt thanks to **Mr. Adithya DSA** from the Data Science Department for generously providing the dataset necessary for this project. His assistance greatly facilitated the data preparation and analysis stages.

Their contributions, encouragement, and expertise have been vital to the progress and outcome of this project, and I am deeply thankful for their support.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 5, Issue 4, May 2025
Impact Factor: 7.67

REFERENCES

[1] Gadotti, A., Rocher, L., Houssiau, F., Creţu, A., & De Montjoye, Y. (2024). Anonymization: The imperfect science of using data while preserving privacy. *Science Advances*, *10*(29). https://doi.org/10.1126/sciadv.adn7053

[2] Vovk, O., Piho, G., & Ross, P. (2021). Anonymization Methods of Structured Health Care Data: A Literature review. *Lecture Notes in Computer Science*, 175–189. https://doi.org/10.1007/978-3-030-78428-7_14

[3] Zuo, Z., Watson, M., Budgen, D., Hall, R., Kennelly, C., & Moubayed, N. A. (2021). Data Anonymization for Pervasive Health Care: Systematic Literature Mapping Study. *JMIR Medical Informatics*, *9*(10), e29871. https://doi.org/10.2196/29871

[4] Subramanian, R. (2022). Applications of differential privacy to healthcare. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4005908

[5] Alnemari, A., Romanowski, C., & Raj, R. (2017). *An adaptive differential privacy algorithm for range queries over healthcare data*. In **2017 IEEE International Conference on Healthcare Informatics (ICHI)** (pp. 274–279). IEEE. https://doi.org/10.1109/ICHI.2017.49

[6] Ferianasari, I. (2024). Patient data protection in the digital era (challenges and legal solutions). *International Journal of Social Service and Research*. https://doi.org/10.46799/ijssr.v4i11.1112

[7] Tupsakhare, P. (2024). The next frontier in healthcare security: Leveraging AI for enhanced data protection. *Progress in Medical Sciences, 8*(1), 1–4. https://doi.org/10.47363/pms/2024(8)e119

[8] Gawankar, S., Nair, S., Pawar, V., Vhatkar, A., & Chavan, P. (2024). *Patient privacy and data security in the era of AI-driven healthcare*. In *Proceedings of the International Conference on Computing Communication Control and Automation*. https://doi.org/10.1109/ICCUBEA61740.2024.10775004

[9] Murdoch, B. (2021). Privacy and artificial intelligence: Challenges for protecting health information in a new era. *BMC Medical Ethics*, *22*(1), 1-9 https://doi.org/10.1186/s12910-021-00687-3

[10] Xu, R., Baracaldo, N., & Joshi, J. (2021). *Privacy-preserving machine learning: Methods, challenges and directions*. arXiv. https://arxiv.org/abs/2108.04417

[11] Agarwal, R. (2023). *Federated learning in privacy-preserving machine learning: Balancing model accuracy and data security*. Journal for Research in Applied Sciences and Biotechnology. https://doi.org/10.55544/jrasb.2.4.31

[12] Mohassel, P., & Zhang, Y. (2017). *SecureML: A system for scalable privacy-preserving machine learning*. In *Proceedings of the IEEE Symposium on Security and Privacy* (pp. 19–38). IEEE. https://doi.org/10.1109/SP.2017.12

[13] Jia, Q., Guo, L., Jin, Z., & Fang, Y. (2018). Preserving model privacy for machine learning in distributed systems. *IEEE Transactions on Parallel and Distributed Systems, 29*(8), 1765–1778. https://doi.org/10.1109/TPDS.2018.2809624

[14] Xu, M., Yu, B., & Wang, F. (2019). Predictive modeling of the hospital readmission risk from patients' claims data using machine learning: A case study on COPD. *Scientific Reports*, *9*(1), 139. https://doi.org/10.1038/s41598-019-39071-y

[15] Échevin, D., Li, Q., & Morin, M. (2017). *Hospital readmission is highly predictable from deep learning*. (3 citations).

[16] Ramírez, J., & Herrera, D. A. (2019). Prediction of diabetic patient readmission using machine learning. *2019 IEEE Colombian Conference on Applications in Computational Intelligence (ColCACI)*, 1–6. https://doi.org/10.1109/ColCACI.2019.8781796

[17] Silva, H., Basso, T., & Moraes, R. L. O. (2017). Privacy and data mining: Evaluating the impact of data anonymization on classification algorithms. *European Dependable Computing Conference*. https://doi.org/10.1109/EDCC.2017.17

[18] Kundu, S., & Suthaharan, S. (2019). Privacy-preserving predictive model using factor analysis for neuroscience applications. *2019 IEEE 5th International Conference on Big Data Security on Cloud (BigDataSecurity), IEEE International Conference on High Performance and Smart Computing (HPSC), and IEEE International Conference on Intelligent Data and Security (IDS)*, 1-8. https://doi.org/10.1109/BigDataSecurity-HPSC-IDS.2019.00023

[19] Raea, S. M., Almotairi, K. M., Alharbi, A. M., Almutairi, G. T., Alhassun, A. M., A Binselm, K. R., Alruqaie, R. I. M., Albalawi, B. M., Alyamani, A. M. A., Alamri, B. J., Alharbi, M. A., Alsulami, A. N. M., Albather, M. H., &

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 5, Issue 4, May 2025
Impact Factor: 7.67

Alfawzan, I. S. A. (2023). Ethical considerations in the use of patient medical records for research. *International Journal of Health Sciences*. https://doi.org/10.53730/ijhs.v7ns1.15415

[20] Garfinkel, S., Abowd, J., & Powazek, S. (2018). Issues encountered deploying differential privacy. *Proceedings of the Workshop on Privacy in the Electronic Society* (WPES@CCS), 1-10. https://doi.org/10.1145/3267323.3268949

[21] Panavas, L., Sarker, A., Di Bartolomeo, S., Sarvghad, A., Dunne, C., & Mahyar, N. (2024). Illuminating the landscape of differential privacy: An interview study on the use of visualization in real-world deployments. *IEEE Transactions on Visualization and Computer Graphics*. https://doi.org/10.1109/TVCG.2024.3427733

[22] Watson, K., & Payne, D. M. (2020). Ethical practice in sharing and mining medical data. *Journal of Information Communication and Ethics in Society*, *19*(1), 1–19. https://doi.org/10.1108/jices-08-2019-0088

[23] Panavas, L., Sarker, A., Di Bartolomeo, S., Sarvghad, A., Dunne, C., & Mahyar, N. (2024). Illuminating the Landscape of Differential Privacy: An interview study on the use of visualization in Real-World deployments. *IEEE Transactions on Visualization and Computer Graphics*, 1–16. https://doi.org/10.1109/tvcg.2024.3427733

[24] Truex, S., & Malan, M. (2024). Privacy in practice: Research challenges in the deployment of privacy-preserving ML. *2024 IEEE 6th International Conference on Trust, Privacy and Security in Intelligent Systems, and Applications (TPS-ISA)*, 157–165. https://doi.org/10.1109/TPS-ISA62245.2024.00027

[25] Drechsler, J., & Bailie, J. (2023). *The complexities of differential privacy for survey data*. Social Science Research Network. https://doi.org/10.2139/ssrn.4950556

[26] Dwork, C., Kohli, N., & Mulligan, D. (2019). Differential privacy in practice: Expose your epsilons! *Journal of Privacy and Confidentiality*, *11*(1). https://doi.org/10.29012/jpc.689