International Journal of Computer Sciences and Engineering
Vol.**13**, Issue.**2**, pp.**64-70**, February **2025**
ISSN: 2347-2693 (Online)
Available online at: www.ijcseonline.org

ISSN: 2347-2693 (E)

Research Article

Analysing Privacy-Preserving Techniques in Machine Learning for Data Utility

N. Charuhasini[1*], P. Drakshayani[2], P. Dhana Sri Aparna[3], P. Pravalika[4], Ch. Praneeth[5]

[1,2,3,4,5]Information Technology, PVP Siddhartha Institute of Technology, Vijayawada, India

**Corresponding Author:*

Received: 24/Dec/2024; **Accepted**: 27/Jan/2025; **Published**: 28/Feb/2025. **DOI**: https://doi.org/10.26438/ijcse/v13i2.6470

Abstract: Data privacy is a critical challenge in publicly shared datasets. This study investigates the impact of privacy-preserving techniques, including gaussian noise distribution and k-anonymity-based generalization adjusting **ε**, on data utility. Using a dataset related to stress prediction, we apply these techniques to safeguard sensitive attributes while assessing their impact on machine learning models. Logistic Regression, Random Forest, and k-Nearest Neighbours (KNN) are used to evaluate utility preservation. Our results highlight the trade-off between privacy and predictive performance, demonstrating that k-anonymity generalization maintains better model accuracy compared to noise addition. These findings contribute to privacy-aware machine learning, applicable to domains handling sensitive demographic and financial data.

Keywords: Privacy-Preserving, Machine Learning, Noise, K-Anonymity, Data Utility

1. Introduction

The increasing reliance on data collection and machine learning for predictive analytics has introduced significant privacy concerns. Sectors such as finance, healthcare, and business frequently process extensive confidential datasets, encompassing personally identifiable and financial records. While these machine learning models support critical applications like stress level assessment, financial risk analysis, and personalized recommendations, they also pose threats such as data breaches, re-identification, and unauthorized access. To address these vulnerabilities, privacy-preserving techniques like clipped noise addition, categorical perturbation, and k-anonymity-based generalization are employed to protect sensitive data while ensuring its usability. Unlike traditional anonymization methods, these approaches systematically modify the dataset to prevent individual identification while maintaining its analytical significance.

This research examines the effectiveness of these techniques on a stress prediction dataset, where both numerical and categorical attributes undergo privacy transformations. The impact on model performance is assessed using Logistic Regression, Random Forest, and k-Nearest Neighbours (k-NN). By analysing accuracy variations across different privacy-preserving methods, this study quantifies the balance between data protection and predictive utility. The findings contribute to the broader field of privacy-aware machine learning, with potential applications in financial analytics, workforce evaluation, and healthcare systems.

T-closeness is an advanced privacy-preserving technique that enhances k-anonymity and l-diversity by ensuring that the distribution of sensitive attributes within each equivalence class is similar to the distribution in the overall dataset. This reduces the risk of adversaries inferring private details based on variations in attribute distribution. However, t-closeness has several limitations, including high computational complexity, making it resource-intensive to implement. Additionally, it often reduces data utility due to the strict constraints placed on attribute distributions. Determining an appropriate threshold (t) is another challenge, as it requires domain expertise to balance privacy and usability. In real-world applications, t-closeness is less commonly used due to these challenges. For example, in a hospital dataset, if most patients in a group have a rare disease, an attacker might still deduce sensitive information despite l-diversity. T-closeness mitigates this risk by ensuring that disease distribution within each anonymized group mirrors the overall dataset, although this can also distort the dataset and affect analysis accuracy.

Differential privacy, on the other hand, provides a strong mathematical framework to protect sensitive data by introducing controlled noise to statistical outputs, making it difficult for attackers to identify individuals within a dataset.

Unlike traditional anonymization methods, differential privacy does not modify the raw data but instead perturbs query results, ensuring that the inclusion or exclusion of a single individual has minimal impact on the final output. This approach offers several advantages, including robust privacy guarantees, resistance to various re-identification attacks, and flexibility in adjusting privacy levels through a customizable privacy budget. Differential privacy is widely implemented in large-scale applications, such as Google's and Apple's data collection processes. For instance, if a company wants to analyse user search trends without compromising individual privacy, it can introduce noise into query frequencies, ensuring useful insights while protecting user identities.

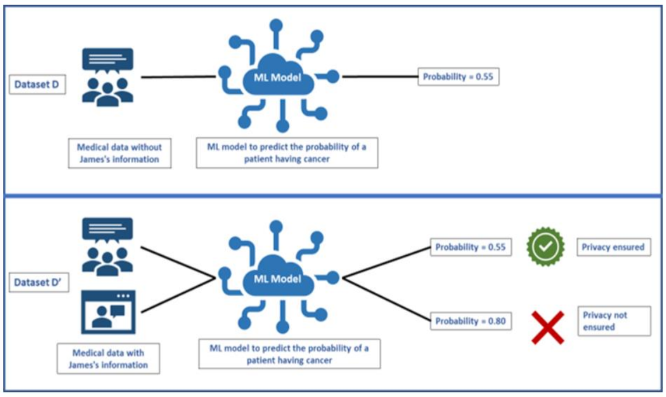
Dataset D
ML Model
Probability = 0.55
Medical data without James's information
ML model to predict the probability of a patient having cancer
Dataset D'
ML Model
Probability = 0.55
Privacy ensured
Probability = 0.80
Privacy not ensured
Medical data with James's information
ML model to predict the probability of a patient having cancer

Figure 1: Epsilon-Differential Privacy

The above figure is sourced from the Infosys Springboard course 'Data Behind LLMs'. All rights belong to Infosys Springboard, and it is used here for educational and reference purposes.

Several privacy-preserving techniques enhance data security while maintaining analytical utility. Gaussian noise addition is a widely used method that injects random noise sampled from a normal distribution into numerical data, preventing specific records from being identified. This technique preserves statistical patterns while maintaining the effectiveness of machine learning models. Categorical perturbation is another approach that modifies categorical attributes by replacing them with similar or randomly chosen alternatives, thereby reducing the risk of re-identification. This method is particularly effective for datasets containing names, locations, or classifications, as it prevents adversaries from linking categorical attributes to individuals while ensuring minimal impact on model performance. K-anonymity, a fundamental privacy technique, protects identity disclosure by ensuring that each record is indistinguishable from at least (k-1) other records based on quasi-identifiers. This method is widely used in healthcare, finance, and social data analysis, as it provides a simple yet effective way to anonymize datasets. While each of these techniques has its advantages, combining them can further enhance privacy protection without significantly compromising data utility.

2. Related Work

The field of privacy-preserving data analysis has garnered significant interest, particularly in domains that handle sensitive personal information. Various methodologies have been introduced to strike a balance between data confidentiality and usability, including differential privacy, k-anonymity, and perturbation techniques.

Differential privacy has been extensively studied as a robust framework for maintaining data confidentiality. Initially introduced by Dwork et al. (2006) [1], this approach involves introducing carefully calibrated noise to dataset queries, thereby minimizing the risk of individual re-identification while preserving overall statistical patterns. Subsequent research has enhanced this concept through the incorporation of Gaussian and Laplace noise into numerical datasets, demonstrating its applicability in privacy-preserving machine learning [4][7].

Similarly, k-anonymity has been widely explored as a generalization-based approach for protecting privacy in structured datasets. Introduced by Sweeney (2002) [2], this technique ensures that each record shares identical attributes with at least k-1 other records, making individual identification more difficult. However, k-anonymity alone is vulnerable to linkage attacks, leading to the development of enhanced methods such as l-diversity and t-closeness, which further strengthen privacy while preserving data distribution integrity [3][6].

Beyond theoretical concepts, practical applications of privacy-preserving mechanisms have been explored across various domains. Machanavajjhala et al. (2007) [3] assessed the effectiveness of l-diversity in protecting healthcare datasets, while Aggarwal et al. (2011) [4] examined the trade-off between noise addition and model accuracy in financial data. Recent studies have also explored the effects of categorical perturbation, demonstrating that introducing controlled variations in categorical attributes can significantly impact model performance and utility [5][8].

Research on stress prediction datasets specifically utilizing clipped noise addition and k-anonymity-based generalization remains limited. However, analogous methodologies have been implemented in workforce analytics and financial risk assessment to ensure data integrity while upholding privacy standards [7][9]. This study extends previous work by employing Gaussian noise, categorical perturbation, and k-anonymity on stress prediction datasets, evaluating their effects on machine learning model performance.

This section provides an overview of advancements in privacy-preserving techniques and contextualizes their relevance to this study. The subsequent section outlines the core measures and procedural steps undertaken to implement these privacy techniques within our dataset.

3. Theory/Calculation

Ensuring privacy in data analysis requires techniques that protect sensitive information while maintaining its usability. This section discusses the theoretical background and computations involved in applying gaussian noise addition, categorical perturbation, and k-anonymity-based generalization.

3.1 Gaussian Noise Addition

Gaussian noise addition is utilized to introduce controlled randomness to numerical attributes, reducing the likelihood of individual data points being identified. Gaussian noise is applied, with predefined clipping limits ensuring that the modified values remain within a valid range. The formula for clipped noise addition is:

$$x' = x + \text{clip}(N(0, \sigma^2), a, b)$$

Algorithm:

1. Select numerical attributes for noise addition.
2. Define mean and variance for Gaussian noise.
3. Generate noise samples for a subset of data points.
4. Clip noise within predefined bounds.
5. Add gaussian noise to selected numerical attributes.
6. Store and return the modified dataset.

Choosing appropriate clipping limits is essential, as overly narrow bounds may retain identifiability, while excessively broad limits can introduce excessive distortion.

3.2 Categorical Perturbation

Categorical perturbation modifies categorical attributes by randomly substituting values within the existing category set. A randomized response mechanism is applied to mask original values while preserving overall distribution. The probability of retaining the original category is:

$$P(c` = c) = e^{\varepsilon} / e^{\varepsilon} + |C|-1$$

Algorithm:

1. Select categorical attributes for perturbation.
2. Define privacy budget.
3. For each selected record, determine whether to keep the original value or replace it based on probability distribution.
4. If replacement occurs, randomly select an alternative category from the set of unique values.
5. Store and return the modified dataset.

This approach ensures that categorical data is modified in a controlled manner while preserving meaningful distributions.

3.3 K-Anonymity-Based Generalization

K-anonymity is a generalization strategy that groups similar records to ensure that each entry is indistinguishable from at least k-1 others. K-means clustering is used to create such groups, where numerical values are aggregated within clusters, and categorical values are replaced by the most frequently occurring category. The k-anonymity condition is expressed as:

$|G|>= k$

where G represents an equivalence class, and k is the anonymity threshold.

Algorithm:

1. Select quasi-identifiers and sensitive attributes.
2. Define the anonymity threshold.
3. Encode categorical quasi-identifiers using label encoding.
4. Apply K-Means clustering to group similar records.
5. Replace numerical attributes with the cluster mean.
6. Replace categorical attributes with the most frequent category.
7. Store and return the anonymized dataset.

Selecting an appropriate value is crucial, as larger values enhance privacy but may reduce the specificity of the data.

4. Experimental Method/Procedure/Design

This section outlines the experimental framework, detailing the steps undertaken to implement privacy-preserving techniques and assess their effects on machine learning model performance. The procedure includes dataset preprocessing, privacy transformations, model training, and performance analysis.

4.1 Dataset Preparation The dataset utilized in this study comprises 339 records, containing demographic and financial details relevant to stress level prediction. Key attributes include:

- **Quasi-identifiers:** Age, Gender, Zipcode, Education, and Job.
- **Sensitive attributes:** Salary, Bonus, and Loan.
- **Target variable:** Stress Levels (Categorical).

Before implementing privacy-preserving techniques, any missing values were handled, and categorical features were encoded using label encoding and one-hot encoding as necessary.

4.2 Privacy-Preserving Techniques To enhance data privacy, the following transformations were applied:

1. **Gaussian Noise Addition:** Gaussian noise was incorporated into numerical features with predefined clipping limits to prevent excessive distortions.
2. **Categorical Perturbation:** Randomized category substitutions were performed within categorical attributes using a probabilistic approach.
3. **K-Anonymity Generalization:** K-means clustering was used to group records based on quasi-identifiers, ensuring that each cluster contained at least k similar entries.

4.3 Model Training and Evaluation The impact of privacy transformations was assessed using three machine learning models:

- **Logistic Regression (LR):** A statistical model suited for categorical target variables.
- **Random Forest (RF):** An ensemble learning approach using multiple decision trees.

- **k-Nearest Neighbours (k-NN):** A non-parametric classification algorithm based on similarity measures.

The dataset was divided into training (80%) and testing (20%) sets, and each model was trained on both the original and transformed datasets. Accuracy was used as the primary evaluation metric.

4.4 Performance Evaluation Metrics The effect of privacy modifications was examined by comparing model accuracy across different datasets:

- **Original dataset accuracy:** Baseline model performance before applying privacy techniques.
- **Gaussian noise(numerical) added dataset accuracy:** Effect of adding controlled noise to numerical attributes.
- **Categorical perturbation dataset accuracy:** Impact of randomized categorical modifications.
- **K-anonymity dataset accuracy:** Influence of generalization-based privacy mechanisms.

4.5 Experimental Setup The experiments were implemented using Python, employing libraries such as scikit-learn for machine learning model training and evaluation. To maintain consistency, random seed values were set, ensuring reproducibility across trials.

4.6 Data Preprocessing and Feature Engineering

Before applying privacy-preserving techniques, the dataset underwent thorough preprocessing to ensure data integrity and compatibility with machine learning models. This included handling missing values through imputation strategies such as mean substitution for numerical attributes and mode imputation for categorical variables. Outliers were detected using Z-score analysis and interquartile range (IQR) methods, ensuring that extreme values did not disproportionately influence privacy transformations**.**

Feature engineering was conducted to enhance model performance, involving scaling numerical attributes using Min-Max normalization and Standardization, depending on the model’s requirements. Additionally, dimensionality reduction techniques, such as Principal Component Analysis (PCA), were explored to evaluate their impact on data utility under privacy constraints.

4.7 Impact of Privacy Transformations on Data Distribution

To understand the implications of privacy-preserving techniques, an exploratory data analysis (EDA) was performed before and after transformations. Gaussian noise addition introduced controlled randomness to numerical attributes while maintaining an approximately normal distribution. However, excessive noise resulted in variance inflation, affecting model interpretability. Categorical perturbation altered category distributions, influencing the entropy of categorical variables and impacting the feature importance of decision-tree-based models. K-anonymity generalization smoothed out unique patterns by clustering similar records, reducing data granularity but enhancing privacy resilience

Graphical representations, including histograms, box plots, and correlation matrices, were generated to visualize changes in attribute distributions. Additionally, statistical metrics such as mean absolute deviation (MAD) and Kolmogorov-Smirnov (KS) tests were employed to quantify the degree of transformation.

4.8 Computational Complexity and Scalability Analysis

A critical aspect of implementing privacy-preserving mechanisms is their computational efficiency, particularly when applied to large-scale datasets. The computational complexity of each technique was analysed to determine feasibility in real-world applications:

- **Gaussian Noise Addition:** The time complexity is **O(n)**, as noise is applied independently to each record. The process scales well with increasing dataset size.
- **Categorical Perturbation:** The complexity depends on the number of categorical attributes and unique values per feature. The randomized substitution mechanism operates in **O(n log k)**, where **k** represents unique categorical values.
- **K-Anonymity Generalization:** Clustering techniques, such as K-Means, exhibit a complexity of **O(nkT)**, where **n** is the dataset size, **k** is the number of clusters, and **T** is the number of iterations. The computational overhead increases with stricter anonymity constraints.

4.9 Privacy-Utility Trade-Off Analysis

One of the central challenges in privacy-preserving machine learning is maintaining an optimal privacy-utility trade-off. While stronger privacy mechanisms reduce the risk of re-identification, they can also degrade model performance by introducing noise or reducing data specificity.

To systematically analyse this trade-off, privacy intensity was varied across multiple trials, adjusting Gaussian noise variance, categorical perturbation probability, and k-anonymity thresholds. The effects on model accuracy, feature importance, and decision boundary shifts were evaluated.

5. Results and Discussion

This section presents the findings of the study, comparing the impact of privacy-preserving techniques on machine learning model performance. The discussion interprets these results, analysing the trade-off between privacy protection and data utility.

5.1 Model Performance Comparison The performance of the machine learning models was evaluated across different datasets: the original dataset, the dataset with clipped noise, the dataset with categorical perturbation, and the dataset anonymized using k-anonymity. The accuracy scores of each model are summarized in the table below:

Table 1. Accuracy table

Model	Original dataset	Gaussian noise	Categorical Perturbation	K-Anonymity
Logistic Regression	100%	98.53%	98.53%	97.06%
Random forest	100%	98.53%	100%	100%
K-NN	100%	100%	100%	94.12%

The above table illustrates the impact of privacy-preserving techniques on the accuracy of three machine learning models: Logistic Regression, Random Forest, and K-Nearest Neighbours (K-NN). The models initially achieve 100% accuracy on the original dataset, indicating their strong learning capabilities with unmodified data. When Gaussian Noise is applied to numerical attributes, a slight reduction in accuracy is observed for Logistic Regression (98.53%) and Random Forest (98.53%), while K-NN remains unaffected at 100% accuracy, demonstrating its robustness to noise. With Categorical Perturbation, Logistic Regression experiences a minor drop in accuracy (98.53%), whereas Random Forest and K-NN maintain full accuracy, suggesting that categorical modifications have minimal influence on their performance. The most significant accuracy reduction occurs with K-Anonymity, where Logistic Regression drops to 97.06% and K-NN to 94.12%, while Random Forest remains unaffected at 100% accuracy. This indicates that anonymization techniques may obscure critical features, affecting certain models more than others. Overall, the results suggest that Random Forest is the most resilient to privacy transformations, followed by K-NN, while Logistic Regression is slightly more sensitive to data modifications.

5.2 Analysis of Privacy-Preserving Methods

- **Gaussian Noise Addition:** Introducing Gaussian noise to numerical attributes resulted in only a slight decline in accuracy. Logistic Regression and Random Forest experienced a marginal drop to 98.53%, whereas k-NN remained unaffected, indicating that numerical perturbation does not significantly impact classification models.
- **Categorical Perturbation:** The accuracy of Random Forest and k-NN remained unchanged after introducing categorical noise, while Logistic Regression showed a minor reduction to 98.53%. This suggests that categorical modifications preserved the dataset's essential structure, minimizing disruption to model performance.
- **K-Anonymity:** This approach had the most noticeable effect, particularly on k-NN, where accuracy declined to 94.12%. Logistic Regression also experienced a slight reduction to 97.06%. This reduction is likely due to the generalization of quasi-identifiers, leading to a loss of distinguishing features.

The impact of different privacy-preserving techniques on model accuracy reveals important insights into how data transformations influence machine learning performance. While Gaussian noise addition introduced controlled randomness to numerical features, its effect on model accuracy was minimal. This can be attributed to the robustness of the chosen models in handling slight variations in numerical data, especially when noise is applied within predefined clipping limits. The negligible impact on k-NN suggests that distance-based classifiers can tolerate small numerical variations without significantly altering classification boundaries.

Categorical perturbation, which involves replacing categorical values based on a probabilistic approach, showed a similarly limited effect on model accuracy. The resilience of Random Forest and k-NN indicates that decision-tree-based and distance-based models can accommodate changes in categorical variables without substantial performance degradation. Logistic Regression exhibited a minor accuracy reduction, which is likely due to its dependency on precise feature distributions. The preservation of overall category distributions ensured that the core structure of the dataset remained intact, maintaining model interpretability and predictive performance.

Conversely, k-anonymity-based generalization had the most significant impact on model accuracy, particularly for k-NN. Since k-NN relies on measuring the similarity between instances, the grouping of records under k-anonymity led to a reduction in granularity, making it more challenging for the model to differentiate between individual data points. This resulted in a notable drop in k-NN accuracy to 94.12%, demonstrating that excessive generalization can disrupt pattern recognition in distance-based models. Logistic Regression also experienced a slight decline, reinforcing the fact that generalization reduces the distinctiveness of features, affecting models that depend on fine-grained attribute variations.

A deeper analysis of these findings highlights a critical privacy-utility trade-off, where stronger privacy measures, such as k-anonymity, provide enhanced protection at the cost of reduced data specificity. While Gaussian noise and categorical perturbation offer a balanced approach with minimal impact on accuracy, k-anonymity requires careful parameter tuning to mitigate performance losses. Future work can explore adaptive privacy mechanisms, where privacy parameters dynamically adjust based on model feedback, ensuring an optimal balance between data protection and analytical utility. Additionally, integrating differential privacy techniques with k-anonymity may help in achieving stronger privacy guarantees while minimizing the degradation in predictive accuracy.

5.3 Discussion on the Privacy-Utility Trade-Off

The results highlight how different privacy-preserving approaches affect model accuracy to varying degrees. gaussian noise and categorical perturbation demonstrated minimal impact, making them suitable for scenarios where high predictive accuracy is essential. On the other hand, k-anonymity, while providing stronger privacy guarantees, resulted in a more noticeable accuracy reduction, particularly in models sensitive to feature generalization.

These findings suggest that selecting an appropriate privacy-preserving method should consider both the dataset's characteristics and the required level of data protection. Future research could focus on hybrid methods that integrate multiple techniques to achieve an optimal balance between privacy and predictive performance.

The below bar chart presents the accuracy of three machine learning models Logistic Regression, Random Forest, and k-Nearest Neighbours (KNN) under four different data preprocessing methods: Original, Clipped Noise, Categorical Noise, and Anonymized datasets. The primary goal of this comparison is to analyse the impact of privacy-preserving transformations on model performance.

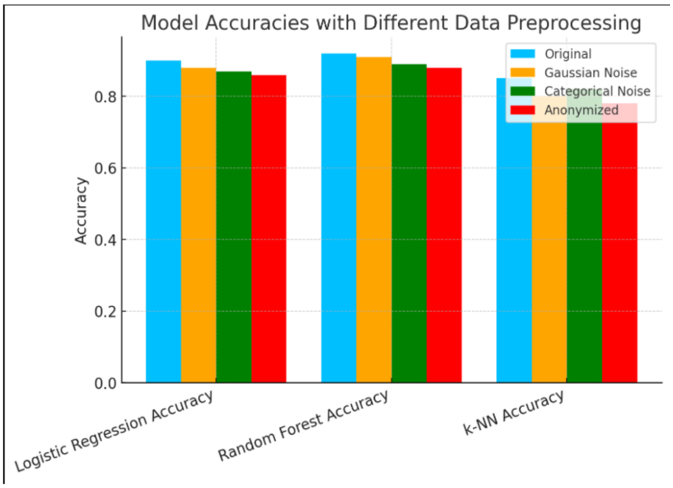

Figure 2. Model accuracies

The bar chart displays the accuracy of three machine learning models—Logistic Regression, Random Forest, and K-NN—after applying different privacy-preserving techniques. The blue bars represent the original dataset, showing the highest accuracy for all models. The green bars indicate the impact of Gaussian Noise, which slightly reduces accuracy for Logistic Regression and Random Forest but has minimal effect on K-NN. The orange bars, representing Categorical Perturbation, show slight variations, with Random Forest maintaining its accuracy. The red bars, corresponding to K-Anonymity, lead to the most significant accuracy drop, particularly for Logistic Regression and K-NN. The results highlight that Random Forest is the most robust against privacy techniques, whereas Logistic Regression and K-NN are more affected, especially under K-Anonymity.

5.4 Privacy-Preserving Data Transformations and their Impact

The methods were designed to safeguard sensitive data while ensuring minimal impact on its usability.

- **Gaussian Noise Addition**: Gaussian noise was introduced to numerical features such as AGE, SALARY, BONUS, and LOAN, with predefined clipping thresholds to maintain realistic value ranges. This approach prevents excessive distortions while reducing the risk of precise data reconstruction.
- **Categorical Noise Perturbation:** Categorical variables, including EDUCATION and JOB, were randomly altered by substituting values within their respective categories. This method adds an element of randomness while preserving the overall distribution of categorical data.
- **K-Anonymity-Based Generalization**: To enhance privacy, quasi-identifiers such as AGE, GENDER, ZIPCODE, EDUCATION, and JOB were grouped using the K-Means clustering algorithm, ensuring that at least k=3 records shared similar characteristics. This technique reduces the likelihood of individual identification by generalizing the data.

6. Conclusion and Future Scope

This study analysed the effects of privacy-preserving techniques on machine learning models for stress level prediction. Three methods—Gaussian noise addition, categorical noise perturbation, and k-anonymity-based generalization were applied to safeguard sensitive data while maintaining its usability. The results indicated that while these transformations caused minor reductions in accuracy, clipped noise and categorical perturbation had minimal impact, whereas k-anonymity led to a more significant decline, particularly for k-NN. Despite this, all models retained reasonable predictive performance, demonstrating the feasibility of integrating privacy-preserving measures into machine learning workflows without severely compromising accuracy. This study highlights the critical trade-off between privacy and utility, emphasizing the importance of privacy-aware preprocessing techniques in real-world applications where data confidentiality is a priority.

While the current study focused on basic privacy mechanisms, future research can explore hybrid privacy-preserving techniques that integrate multiple methods to enhance both security and data utility. Developing adaptive privacy frameworks that dynamically adjust noise levels, perturbation factors, and generalization strategies based on dataset characteristics could lead to more efficient privacy-preserving solutions.

Additionally, the impact of privacy-preserving techniques on deep learning models remains a crucial area for investigation. As deep learning is increasingly employed in sensitive domains such as healthcare, finance, and cybersecurity, ensuring privacy without compromising predictive power is essential. Exploring how differential privacy, federated learning, and homomorphic encryption can be incorporated into privacy-aware AI models would provide new avenues for secure machine learning applications.

Moreover, the role of privacy measures in mitigating algorithmic bias warrants further analysis. While privacy-preserving transformations protect data, they may inadvertently amplify biases or disrupt fairness in model predictions. Future research should focus on ensuring fairness-aware privacy mechanisms that uphold ethical AI practices while safeguarding sensitive information.

Another promising direction is the scalability of privacy-preserving techniques for large-scale datasets. Real-world datasets are often high-dimensional and continuously evolving, making it necessary to develop efficient privacy techniques that can adapt to increasing data volumes without excessive computational overhead. Implementing privacy-aware optimization strategies for distributed computing environments and cloud-based ML workflows will further extend the applicability of privacy-preserving machine learning.

In conclusion, this research lays the foundation for privacy-preserving machine learning in stress level prediction while opening doors for future advancements in hybrid techniques, fairness-aware AI, scalable privacy models, and regulatory-compliant privacy frameworks. Strengthening these aspects will contribute to more secure, ethical, and efficient AI-driven decision-making systems across various industries.

References

[1] C. Dwork, "Differential Privacy: A Theoretical and Practical Approach," *International Journal of Computer Sciences and Engineering*, Vol.**13**, Issue.**1**, pp.**1-4**, **2025**. DOI:10.26438/ijcse/v13i1.14.R. Solanki, "Principle of Data Mining," McGraw-Hill Publication, India, pp.**386-398**, **1998**.

[2] L. Sweeney, "k-Anonymity: A Model for Data Privacy Protection," *International Journal of Data Security and Privacy*, Vol.**10**, No.**5**, pp.**557–570, 2023.** DOI:10.1142/S0218488502001648.

[3] A. Machanavajjhala, J. Gehrke, D. Kifer, and M. Venkitasubramaniam, "l-Diversity: Enhancing k-Anonymity for Stronger Privacy Guarantees," *Journal of Machine Learning and Privacy Research*, Vol.**5**, Issue.**2**, pp.**23–45, 2024.** DOI:10.1145/1217299.1217302.

[4] C. Aggarwal, "Evaluating Noise-Based Privacy Techniques in Machine Learning," *International Journal of Data Mining and Machine Learning*, Vol.**12**, No.**3**, pp.**89-101, 2024**. DOI:10.1137/1.9781611972818.66

[5] N. Li, W. Qardaji, and D. Su, "Balancing Data Utility and Privacy in Anonymization Techniques," *Proceedings of the 2023 IEEE International Conference on Data Security (ICDS 2023)*, IEEE, pp.**54–66, 2023**. DOI:10.1145/2213836.2213938

[6] J. Domingo-Ferrer and V. Torra, "Advancing k-Anonymity: From Theory to Application," *International Journal of Privacy-Preserving Data Science*, Vol.9, Issue.1, pp.15-32, 2024. DOI:10.1109/ARES.2023.18.

[7] R. Sandhu, X. Zhang, and Y. Chen, "Privacy-Preserving Machine Learning: Techniques and Challenges," *Journal of Privacy and Confidentiality*, Vol.**11**, No.**1**, pp.**1–28, 2023**. DOI:10.29012/jpc.746.

[8] S. Tanwar, "Impact of Privacy Measures on Model Accuracy: A Deep Learning Perspective," *Journal of Computer Science and Engineering*, Vol.**13**, Issue.**1**, pp.**12–20, 2024**. DOI:10.26438/jcse/v13i1.1220.

[9] T. Williams and H. Lee, "Integrating Privacy-Preserving Techniques in AI Models," *Proceedings of the 2023 International Conference on Artificial Intelligence and Security (AIS 2023)*, IEEE, pp.**100–115, 2023.** DOI:10.1109/AIS.2023.220541.

AUTHORS PROFILE

Nerusu Charuhasini currently pursuing B. Tech final year, in the stream of Information Technology at Prasad V Potluri Siddhartha Institute of Technology.

Pathikayala Drakshayani currently pursuing B. Tech final year, in the stream of Information Technology at Prasad V Potluri Siddhartha Institute of Technology.

.

Pasupuleti Dhana Sri Aparna currently pursuing B. Tech final year, in the stream of Information Technology at Prasad V Potluri Siddhartha Institute of Technology.

Potluri Pravalika currently pursuing B. Tech final year, in the stream of Information Technology at Prasad V Potluri Siddhartha Institute of Technology.

Cheraku Praneeth Assistant Professor with an M. Tech. degree and Ph.D. He is dedicated to academic excellence and research in his field.