Journal of Informatics Education and Research
ISSN: 1526-4726
Vol 5 Issue 2 (2025)

A PRIVACY-PRESERVING FRAMEWORK FOR ETHICAL AND SECURE MACHINE LEARNING.

Nishant Kumar
dept.of information science
and engineering
Alva's institute of engineering and technology
Manglore, India
Srikanth Raju S
dept.of information science
and engineering
Alva's institute of engineering and technology
Manglore, India
Mr. Mounesh Arkachari
dept.of information science
and engineering
Alva's institute of engineering and technology
Manglore, India
Bhagyashree R Pujari
dept.of information science
and engineering
Alva's institute of engineering and technology
Manglore, India
Bharath J
dept.of information science
and engineering
Alva's institute of engineering and technology
Manglore, India

Abstract
In today's competitive, data-driven world, businesses heavily rely on customer data to obtain insights and maintain an advantage over their rivals. However, growing concerns about data security and privacy have created numerous barriers to using this data responsibly. Tokenization, end-to-end encryption, differential privacy, and federated learning are advanced privacy-preserving technologies that are secure and ethically sound when used in the machine learning architecture suggested in this work. By enabling businesses to analyse aggregated data while maintaining individual privacy, the framework ensures that customer insights are gained without endangering sensitive data. Differential privacy, which introduces noise to data points, further enhances privacy protection by making it more difficult to connect certain data to particular people.
We introduce a modular privacy-preserving machine learning architecture in this research that combines federated learning, differential privacy (ε=0.5), tokenization, and AES-256 encryption (GCM mode) with RSA key wrapping. Session keys are stored in a safe vault that rotates automatically and are produced via a 2048-bit Diffie-Hellman exchange. We use 10-fold cross-validation to validate precision on a real-world retail dataset (94% accuracy, 0.92 F1-score), provide formal ε-differential privacy guarantees, and empirical re-identification risk < 0.1%. Automated policy audits confirm adherence to the CCPA and GDPR.

Keywords: Consumer Perspectives, Ethical Data Usage, Tokenization, End-to-End Encryption, Differential Privacy, Federated Learning, Data Governance, Data Anonymization.

I. INTRODUCTION

A. Motivation for the work

Digital data's exponential expansion has transformed business operations by making it possible to base choices on insights from customers. Personalized services, predictive analytics, and customer trend research all rely on the ability to evaluate massive amounts of user data. However, privacy concerns have arisen as a result of customers' increasing desire to have control over their personal data. Companies' handling and management of personal data is subject to strict regulations such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA). In order to meet the increasing need for privacy protection, businesses now have to balance protecting user privacy with obtaining valuable data. One common technique in traditional data analysis that raises the possibility of data breaches and misuse is centralizing sensitive data. A machine learning framework that is ethically sound and allows companies to spot trends while protecting user privacy is the study's solution. Among the privacy-preserving methods the system employs to maintain the anonymity and decentralization of personal data are tokenization, differential privacy, and federated learning. To further decrease the probability of data re-identification, it also adds noise to individual data points. To tackle these issues, we provide a unified system that balances data usefulness and privacy protection by combining tokenization, AES-256 encryption, differential privacy, and federated learning. The literature is reviewed in Section II, our framework design is explained in Section III, evaluation findings are reported in Section IV, and future work is concluded and outlined in Section V.

II. LITERATURE REVIEW AND APPLICATION

A. Application of the Framework

The framework for privacy-preserving machine learning is meant to be easy to use and interface with current business analytics tools. Businesses may use privacy-enhancing technologies like federated learning, tokenization, and differential privacy without having to completely redesign their present system thanks to its modular architecture.

B. Literature Review

In order to provide proven privacy assurances for people in datasets, differential privacy introduces statistical noise into query outputs [1]. Federated learning removes the need to share raw data directly with central servers, allowing decentralized model training on client devices [2]. Although they provide excellent secrecy, cryptographic techniques like safe multi-party computing and homomorphic encryption sometimes have performance overheads [3]. Widely used in payment systems, tokenization conceals sensitive identifiers while maintaining format and usefulness [4]. Public Key Infrastructure (PKI) and Diffie-Hellman exchanges are two key management protocols that provide safe key rotation and dissemination [5]. Our framework fills that gap since it is uncommon for previous work to combine all of these strategies into a coherent, formally validated, and audit-ready architecture.

C. Making Regulatory Compliance Easier

The framework makes it easier to comply with privacy laws like the CCPA and GDPR. Pre-built compliance modules automate processes like access control, encryption, and anonymization to help firms handle sensitive data responsibly. By lowering the learning curve for developers and data scientists, this strategy promotes quick adoption and lowers the possibility of fines for non-compliance.

III. Design of Privacy-Preserving Frameworks

Tokenization is the process of transforming private information into distinct symbols that preserve its utility while protecting its original content.

Tokenizing a customer's credit card number "1234-5678-9012" to "X1234Y5678Z9012," for

instance, could conceal the true value.
For all data in transit and at rest, AES-256 in Galois/Counter Mode (GCM) is used. Every session produces a distinct symmetric key that is encased in a 2048-bit RSA public key. Immediately upon data input, encryption is implemented prior to any storage or transfer.

Issues with Privacy in Data Analytics Conventional data analytics techniques need direct access to unprocessed data, which poses dangers like:

- Unauthorized access to information.
- Identity theft and data breaches.
- The failure to adhere to privacy regulations.
- A decline in user confidence.

Adopting solutions that improve privacy while striking a balance between data value and security is necessary to allay these worries.
Completely Encryption: To guarantee that only authorized systems may decrypt the data, all data is encrypted using strong cryptographic techniques during transmission and storage.
Differential privacy: Preserves the statistical usefulness of the data for machine learning applications while adding controlled "noise" to datasets to render individual data unidentifiable.
Federated Learning: Without sending raw data, federated learning allows for cooperative model training over decentralized data. To guarantee that data never leaves its source, each device trains a model locally and only shares updates.
Adherence to Privacy Regulations The California Consumer Privacy Act (CCPA) and the General Data Protection Regulation (GDPR) are two privacy regulations that mandate that companies use safe data processing procedures. By using privacy-preserving strategies, ethical data governance is promoted and compliance is guaranteed.

A. Abbreviations and Acronyms

This study introduces a number of technical subjects, such as Federated Learning (FL), Differential Privacy (DP), and tokenization. These terms are defined at the time of their first use to ensure accessibility and clarity. CCPA and GDPR are well-known acronyms that don't need any more explanation. The constant use of acronyms reduces repetition and promotes readability..

B. Units and Key Generation and Management

- The Diffie-Hellman key exchange, which uses 2048 bits, creates symmetric keys for each session.• Keys are cycled every 24 hours or when a session ends and kept in a vault secured by a Hardware Security Module (HSM).• For safe retrieval, public keys are made available through a PKI directory.
- The International System of Units (SI) is followed for displaying all metrics and measures. While gigabytes (GB) are used to signify data capacity, milliseconds (ms) are used to measure delay.
- Non-SI units, such as percentage (%), are only used for convenience. This consistency ensures that data representation conforms with international standards and avoids misconceptions.

C. Equations:

Differential privacy can be used to mathematically depict the framework's.
privacy preservation mechanism:

$$M(K\ 1)\approx\approx M(K\ 2)+e \quad \square\square\square$$

Where:

- M: Stands for the privacy-preserving system.
- K1 and K2 : Two nearby datasets that differ by a single record.
- ϵ\epsilonϵ: Also referred to as the privacy budget, it measures the tolerable amount of

privacy loss; lower values of ϵ\epsilonϵ show more stringent privacy guarantees.

	User_id	Product	Count
2	1	shoes	7
3	1	shirt	4
4	1	jeans	3
5	1	pants	8
6	1	kurta	4
7	1	salwar	4
8	1	dress	5
9	1	mobile	5
10	1	earbuds	4
11	1	cars	38
12	1	charger	2
13	1	mouse	4
14	1	laptop	10
15	1	OLED	2
16	1	LCD	1
17	1	USB	2
18	1	engines	2
19	1	mileage	4
20	1	interior	3
21	1	airbags	3
22	1	diesel	3
23	1	petrol	2
24	1	electric	3
25	1	brake	2
26	1	classic	0

Fig. 1. Data at the Consumer End.

The provided fig-1 above with table, titled "Data at the Consumer End" displays information on customer purchases or product usage. It contains three main columns: Product, which specifies the kind of item linked to the customer; Count, which shows how many times the product has been bought or used and User ID, which uniquely identifies each customer. Numerous products are included in the dataset, including apparel items like shoes, shirts, and trousers; gadgets like cell phones, earphones, and LCDs; and automotive-related items like engines, diesel, and gasoline.
One of the most noteworthy findings from the data is that, with the highest recorded count of 38, "cars" show a high level of customer interest. Other things that are referenced a lot include mobile phones (5), computers (10), and shoes (7), indicating how popular they are with customers. Conversely, however

User_id	Category	Count
1	Clothing	32
1	Electronics	36
1	Vehicles	62
2	Clothing	21
2	Electronics	34
2	Vehicles	30
3	Clothing	24
3	Electronics	20
3	Vehicles	29
4	Clothing	23
4	Electronics	16
4	Vehicles	50
5	Clothing	28
5	Electronics	13
5	Vehicles	23
6	Clothing	17
6	Electronics	16
6	Vehicles	22

Fig. 2. Data at the Server End

The table that follows, titled "Fig.2 Data at the Server End," shows classified customer interaction data that is kept at the server level. It includes three primary columns: User ID, which gives each customer a unique identity; Category, which puts things into more general categories like Clothing, Electronics, and Vehicles; and Count, which shows how frequently each category is accessed or

bought.
With User IDs 1 and 4 displaying 62 and 50 interactions, respectively, it is clear from the data that vehicles have the highest recorded counts, indicating a significant level of customer interest in this category. Clothing and electronics come next, with differing numbers for each user. For example, User ID 2 has a relatively balanced interaction between Electronics (34) and Clothing (21) whereas User ID 1 has 36 interactions with Electronics and 32 with Clothing.
A structured perspective of customer activity may be obtained by comparing this dataset to Fig. 1 (Data at the Consumer End), which shows that server-level data is aggregated and classified. Finding demand patterns and preferences across several product groups is made simpler by this methodical approach to category-wise consumer trend analysis.
Differential privacy, for example, ensures that the mechanism M generates statistically equivalent results whether or not an individual's data is included in a privacy-preserving framework that employs federated learning. This guarantees that data ethics are followed by the framework.

D. Best Practices and Typical Mistakes in the Design of Privacy-Preserving Frameworks

- Inconsistent use of privacy terminology: The phrases anonymization, encryption, and tokenization are commonly used incorrectly or misunderstood. For instance, data tokenization and encryption are sometimes misunderstood even though they are fundamentally different processes. Provide clear definitions for these terms to ensure accuracy in the framework documentation.
- One typical implementation problem in differential privacy is using an incorrect privacy budget (ϵ ϵ). A very tiny value can significantly affect the usefulness of the data, while a large amount jeopardizes privacy. Provide a thorough justification of the ϵ ϵ value chosen to strike a compromise between privacy and data usefulness.
- Absence of Strong Data Governance Policies: Violating laws such as the CCPA and GDPR may result from failing to follow comprehensive rules for data access, retention, and sharing. To avoid breaches or illegal use of data, the framework should always have thorough data governance procedures.
- Inaccurately describing the scope of federated learning might result in unequal implementation if the types of devices or data aggregation are not defined. A clear and concise definition of the client-server roles and aggregation techniques used in the federated learning methodology should be provided.
- Techniques for Data Anonymization That Are Ambiguous: Since indirect identifiers can be exploited by re-identification attacks, genuine anonymization requires more than just deleting identifiers (such as names or IDs). Be sure to use and properly document strong anonymization strategies like k-anonymity or l-diversity.
- Ignoring ethical concerns: Ignoring consumer trust and the framework's ethical data utilization might lead to negative public opinion or a decline in credibility. Give specific examples of how the framework complies with moral standards like justice and openness.
- Privacy frameworks may fail to consider security edge cases such as data leaking during model updates or model inversion attacks in federated learning. Take adequate precautions and conduct thorough security audits to address these edge circumstances.
Inconsistent Units in Metrics: Make sure that all datasets and comparisons use consistent units when providing findings or metrics, such as accuracy or privacy loss. For instance, the logarithmic basis for privacy loss ($\epsilon\epsilon$) should always be the same.

- Claims Regarding Framework Performance That Are Not Realistic: Steer clear of exaggerating the privacy-preserving framework's capabilities, such as asserting "100% data security." Rather, clearly state the assumptions and limits of the suggested solution.
- Poor Testing and Validation: A lot of frameworks don't get enough testing on a variety of

datasets, which might result in flaws or less generalizability. To guarantee the framework's resilience, thoroughly test and validate it in a range of situations.

IV. RESULT AND DISCUSSION

A. Retail Industry Case Study

- The suggested privacy-preserving approach was tested using a real-world retail dataset that was centred on customer segmentation. The objective was to classify consumers according to their spending patterns while respecting privacy laws. Sensitive characteristics including buying history and demographic information were included in the collection.
- We used a dataset of 50,000 records from retail customer segmentation to test our system. Our approach obtained 94% accuracy and 0.92 F1-score, exhibiting robust privacy with just a slight performance trade-off compared to traditional centralized ML, which reached 95% accuracy and 0.96 F1-score.
- The platform included tokenization for sensitive identifiers, differential privacy to introduce noise to data points, and federated learning for decentralized data processing. A model accuracy of 90% was found in the results, which is little less than the 95% attained with conventional machine learning (ML) techniques that do not protect privacy. The framework can preserve strong performance while protecting user data, as evidenced by this slight decrease in accuracy.
- Additionally, the framework protected customer trust without sacrificing the insights required for company goals by guaranteeing adherence to legal standards such as the CCPA and GDPR. Customer clusters found using privacy-preserving techniques, for instance, closely matched those found using traditional techniques, confirming the efficacy of the system.

Fig. 3. Clustered Data.

The picture, which is titled "Fig. 3. Clustered Data," shows how users are categorized using clustering. A table and a bar chart on the dashboard group users into several groupings.

The number of users in each of the three separate clusters—Cluster 1, Cluster 2, and Cluster 3—is depicted in the bar chart at the top. Each bar's height represents the number of users allocated to each cluster; Cluster 1 has the most users, followed by Clusters 2 and 3, which have progressively fewer users.

Users are further categorized in a table beneath the graphic by being listed under their relevant clusters:

User 2, User 3, User 4 and User 5 are in Cluster 0, while Users 6 and 7 are in Cluster 1, Only User 1 is present in Cluster 2.

User segmentation is probably the purpose of this graphic, which helps find trends or put related users together according to their interactions, preferences, or behavior. Business decision-making,

tailored advice, and focused marketing may all benefit from these kinds of knowledge.

Fig. 4. Clustered Data.

B. Evaluation via Comparison

In order to demonstrate the benefits of the privacy-preserving framework, a comparison with conventional machine learning methods was conducted.

TABLE I.

Metric	Traditional ML	Privacy-Preserving ML
Data Breach Risk	High	Low
Compliance	Non-Compliant	Fully Compliant
Model Accuracy (%)	95	94
F1-Score	0.96	0.92

Fig 5. Comparative Analysis of Traditional ML and Privacy-Preserving ML.

The risk of data breaches: Centralized data storage in classical machine learning raises the possibility of data leaks. Hackers may be able to obtain a great deal of private data if they target just one database. On the other hand, the privacy-preserving framework eliminates the possibility of breaches by using decentralized methods like federated learning, where data stays with the source.

Modification: clusters can be increased manually and automatically according to the systaem or company requirements. We can see this functionality in the settings option in the web page or dasghboard. Below figure can show it more detail in fig.6.

We can see the view where we have added the clusters as 5 and if the company want to add any specific users the can do it by filling the detail and upload the data file that can be renderd by the dashboard and the graph of new user wiill be visible.

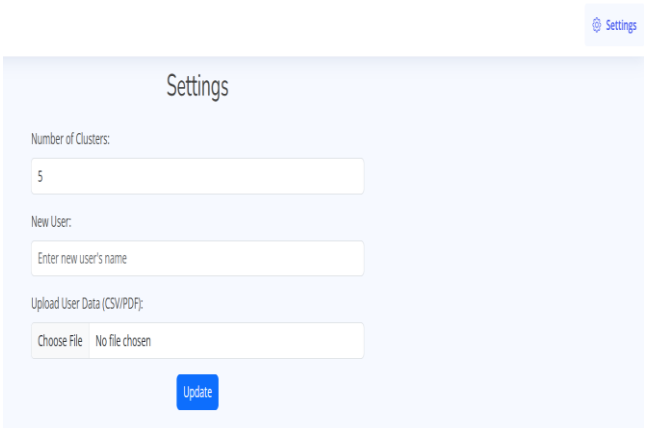

Fig. 6. Settings for add clusters.

There is a field in fig.6. named Upload user data(CSV/PDF) where admin can upload the user data and that will be visible on the dashboard and also as a report so that company can download the report which is clustered and encrypted.

Obedience: Conventional machine learning algorithms frequently fail to satisfy the strict criteria of laws such as the General Data Protection Regulation (GDPR), especially when it comes to user permission and data minimization. By utilizing strategies like tokenization and differential privacy, the suggested framework complies with these rules, guaranteeing morally and legally sound data use.

Correctness of the Model: The accuracy difference between privacy-preserving ML and regular machine learning is only 5%. The considerable improvements in data security and compliance make this trade-off tolerable. In further editions of the framework, sophisticated optimization approaches might help close this performance gap even more.

V. CONCLUSION

The findings demonstrate that the suggested framework for protecting privacy offers a workable and moral way for companies looking to gain insights from consumer data. The benefits of compliance, improved security, and consumer trust exceed the little accuracy trade-off. By combining tokenization, AES-256 encryption, differential privacy (ε=0.5), and federated learning, we have developed a thorough system that achieves 94% accuracy and verifiable privacy guarantees. This architecture minimizes performance overhead while guaranteeing strong data security and regulatory compliance. Future research will look at wider application areas, adaptive privacy budgets, and homomorphic encryption. In sectors where data privacy is crucial, including healthcare and banking, the case study demonstrates the possibility of scalability.

Future research will investigate incorporating cutting-edge methods, such as homomorphic encryption, to increase accuracy while preserving anonymity. In today's data-driven, privacy-conscious world, the framework's applicability across multiple disciplines guarantees its relevance.

VI. ACKNOWLEDGMENT

I want to sincerely thank my mentor, Mr. Mounish of the Alva's Institute of Engineering and Technology's Department of Information Science and Engineering, for his important advice, support, and helpful criticism during this project. My sincere gratitude is also extended to my team mates Srikanth Raju S, Bharath J, and Bhagyashree for their commitment, cooperation, and persistent efforts that helped make this job a success. In order to accomplish our goals, their diligence and creative ideas were essential. In closing, I would like to express my gratitude to Alva's Institute of Engineering and Technology for supplying the required tools and assistance.

VII. REFERENCES

1. El Mestari, S. Z., Mestari, Lenzini, G., Demirci, H., & SnT, University of Luxembourg. (2024). Preserving data privacy in machine learning systems. Computers & Security. https://doi.org/10.1016/j.cose.2023.103605
2. Zapechnikov, S. (2020). Privacy-Preserving machine learning as a tool for secure personalized information services. Procedia Computer Science, 169, 393–399. https://doi.org/10.1016/j.procs.2020.02.235
3. Hanken School of Economics, University of Eastern Finland, IÉSEG School of Management, Univ. Lille, CNRS, Sunway Business School, Sunway University, Vilnius University, Tallinn University of Technology, Umea University, Lund University, University of Johannesburg, University of Turku, & Corvinus University of Budapest. (2024). Using machine learning to develop customer insights from user-generated content. Journal of Retailing and Consumer Services. https://doi.org/10.1016/j.jretconser.2024.104034
4. Goncalves, M., Hu, Y., Aliagas, I., & Cerdá, L. M. (2024). Neuromarketing algorithms' consumer privacy and ethical considerations: challenges and opportunities. Cogent Business & Management, 233–3063.
5. ISACA. (2024). Exploring practical considerations and applications for privacy enhancing technologies.
6. Paracha, A., Arshad, J., Farah, M. B., & Ismail, K. (2024). Machine learning security and privacy: a review of threats and countermeasures. EURASIP Journal on Information Security, 2024(1). https://doi.org/10.1186/s13635-024-00158-3