

A Dual-Layer Privacy-Preserving Federated Learning Framework

Wenxuan Huang^[0000-0002-2613-2672], Thanassis Tiropanis^[1234-5678-9012], and
George Konstantinidis^[0000-0002-3962-9303]

Electronics and Computer Science, University of Southampton
{wh1g19, t.tiropanis, g.konstantinidis}@soton.ac.uk

Abstract. With the exponential growth of personal data use for machine learning models, significant privacy challenges arise. Anonymisation and federated learning can protect privacy-sensitive data at the cost of accuracy but there is lack of research on hybrid approaches. This paper uses federated learning and traditional centralised machine learning to evaluate the effectiveness of different anonymization strategies in environments with independent and identically distributed data. It considers the two layers of data collection (layer one) and model training (layer two) on three scenarios: (i) local data collection and local anonymisation for federated model training, (ii) central data collection before anonymisation for centralised model training, and (iii) central aggregation of locally anonymised data for centralised model training. Our assessment shows that the performance of the models generally decreases with increasing anonymity constraints, but the extent of decrease varies across different scenarios. In addition, we propose a dual-layer federated learning framework that applies differential privacy to ensure privacy during both data collection and model training stages. Evaluation on real-world datasets demonstrates that our framework achieves both acceptable data anonymization and model accuracy.

Keywords: Privacy preservation · Machine Learning · Federated learning · Anonymisation.

1 Introduction

In the era of big data, the amount of data generated by humans is exponentially increasing. A significant portion of this data is inherently personal and includes various types of patient data collected and stored in electronic health records within the electronic health systems. This data encompasses laboratory test results, demographic information, age and weight statistics, as well as medication information[12]. Similarly, in social networks, data such as user names, addresses, email addresses, personal photos, and notes are collected[21]. This data can be used for numerous scientific or commercial purposes, such as data-driven research and product development, relying on the analysis of personal information to generate knowledge-based decisions or provide personalised services. However,

due to the presence of personal information in the data, it may face threats from inference attacks, where attackers can infer sensitive or private information that has not been explicitly disclosed by utilising patterns, correlations, or statistical characteristics present in the available data[20][31]. Therefore, privacy protection is crucial in the publication and utilisation of personal data. The "privacy by design" paradigm[14] emphasises minimising the use of sensitive information. The General Data Protection Regulation (GDPR)[25] introduced by the European Union provides a strict and mandatory framework for safeguarding personal and sensitive information.

To address the threats posed by privacy attacks, anonymization techniques and federated learning are widely employed to protect the publication and usage of privacy-sensitive data[2][18]. In the data collection phase, privacy-preserving data publishing (PPDP) offers a set of models, tools, and methods to mitigate privacy threats when releasing data. Data owners protect users' private information by applying anonymization techniques, including generalisation, suppression, microaggregation, to prevent inference attacks while preserving the utility of the anonymised data. Moreover, in the era of big data, the application of machine learning is becoming increasingly widespread. With the explosive growth of data and the rapid generation of information, traditional methods of data processing and analysis are no longer able to meet the demands of mining and insights from massive data. The development of machine learning technology enables people to harness the valuable information hidden in big data and extract profound insights and patterns, providing more accurate and efficient solutions. However, considering most real-world scenarios, personal data is often scattered across data islands, such as different healthcare institutions or banking systems. However, most traditional machine learning algorithms operate in a centralised manner, requiring data aggregation on data servers. This introduces a single point of failure and significant risks of data breaches, leading to a lack of trust among end users and challenges in complying with GDPR requirements. To overcome these challenges, Google researchers introduced federated learning as a promising solution in 2016, which has gained attention from both industry and academia. Unlike traditional machine learning, federated learning is a framework that enables machine learning algorithms to be implemented in a decentralised collaborative learning setting. In federated learning, models are executed on multiple local datasets stored on various local nodes, such as smartphones, tablets, personal computers, and Internet of Things (IoT) devices[15]. federated learning leverages a central server to coordinate the training process or utilises underlying peer-to-peer network infrastructure to aggregate training results and compute a global model[2]. This allows local nodes to collaboratively train a shared machine learning model, exchanging only the trained parameters (e.g., weights and biases of deep neural networks) periodically, without the need to centrally collect and process training data on a central data server. As a result, federated learning possesses a natural advantage in preserving data privacy. Furthermore, the parameter update and aggregation processes between local nodes and the

central coordinating server can be enhanced with privacy protection techniques, such as differential privacy, to further strengthen data privacy protection[30][9].

However, despite the individual achievements of anonymization and federated learning in privacy protection during different stages of data processing, combining both techniques into a comprehensive privacy protection framework poses challenges. Previous research introducing anonymization algorithms often overlooked the relationship between anonymization and machine learning, partially due to the distinct origins of these two methods. It has been noted that information loss resulting from the generalisation and suppression algorithms in anonymization methods may lead to performance degradation in machine learning models[27][1][4]. Additionally, as a distributed machine learning framework, federated learning distributes data across local nodes instead of aggregating it on a central server. Currently, there is no research that discusses the relationship between anonymization in distributed data collection and the performance of machine learning models. Intuitively, With an increasing number of data holders, each holder possesses a smaller amount of data, resulting in lower levels of anonymity. To achieve anonymity constraints equivalent to those in centralised data sets, the data sets used in federated learning require a greater degree of anonymization transformation, which may further compromise model performance. Therefore, this study first investigates the impact of general anonymization strategies on three different training modes (federated learning, centralised machine learning, hybrid mode), addressing a gap in the existing literature. Subsequently, we propose a novel dual-layer federated learning framework that achieves privacy protection in both the data collection and data processing stages while maintaining acceptable model performance.

The main contributions of this paper are as follows:

- A dual-layer privacy-preserving federated learning framework to explore privacy protection in both the data collection and model training stages while maintaining acceptable model performance.
- Comparison of anonymisation techniques between federated learning on distributed data sets and centralised machine learning on centrally collected data sets using a real-world data set.
- Assessing the impact of different privacy-preserving anonymisation strategies on 3 training scenarios: anonymisation on federated learning, anonymisation on centralised learning, and centralised learning on datasets anonymised before aggregation.
- Evaluation of the proposed framework on a real-world data set, demonstrating improvements in both data anonymity and model performance.

Through these contributions, we provide a comprehensive understanding of the impact of anonymization strategies on federated learning, filling a gap in the field of privacy protection. Moreover, we introduce an innovative federated learning framework that simultaneously enhances data privacy and improves model performance in big data applications. This work is of significant importance in advancing privacy protection and the application of big data, providing valuable guidance for future research and practical implementations.

2 Background

2.1 Challenges in Data Publishing

As more and more personal information is used for data-driven research or product development, the protection of private personal information is becoming increasingly important. PPDP is a process of sharing data while protecting individual privacy. It involves techniques that aim to prevent unauthorised access to sensitive data and protect the anonymity of individuals in the dataset. It is particularly important in industries such as healthcare, finance, and government, where sensitive data needs to be shared for research or analysis purposes while maintaining the privacy of individuals. Table 1 shows the different attributes of the data in the PPDP.

Type	Description
Identifier	Attributes in the data that are used to uniquely identify individuals' identities.
Quasi-identifier	Attributes Combinations of attributes in a dataset that are linked with an Identifying Attribute.
Sensitive Attribute	Attributes in the data that are related to individuals' sensitive information.
Non-sensitive Attribute	Attributes in the data that are not sensitive attributes, identifying attributes, or quasi-identifier attributes.

Table 1: The types of attributes in privacy-preserving data publishing.

In the publishing of personal data, there are three types of privacy threats[8]:

- **Identity disclosure:** An attacker can correctly associate an individual with a personal record in a published dataset.
- **Attribute disclosure:** Attackers can obtain individuals' sensitive information through inference attacks. This type of threat is more likely to occur in datasets with low anonymity.
- **Membership disclosure:** Attackers can infer with a high probability whether an individual's record exists or does not exist in a published dataset[29][23].

2.2 Privacy Paradigms

To protect user privacy in the published dataset, data publishers can apply anonymization techniques to enhance the dataset's resilience against attacks. The aim of anonymisation is to ensure that a person's data record can no longer be traced explicitly back to that particular person. To this purpose, various complementary privacy paradigms have been defined:

k-Anonymity K-anonymity[28] is a privacy protection concept that requires each record in a dataset to be indistinguishable from at least K-1 other records in terms of their attributes, thereby hiding the specific identity information of individuals. K-anonymity is achieved by generalising or suppressing attributes, which increases the similarity between records and ensures anonymity. Specifically, for a dataset D with attribute set A, if for every record d in the dataset, there exist at least K-1 other records d' such that they have the same values for the attributes in set A, then the dataset D satisfies K-anonymity.

L-Diversity L-diversity[19] is a measure of the richness of information in an anatomised dataset. It quantifies the number of different attribute values within each equivalence class in the dataset. The goal of L-diversity is to increase the diversity of attribute values in the economised dataset, thereby enhancing its utility. A higher L-diversity value indicates a greater diversity of attribute values in the dataset.

L-Diversity is measured by:

$$L - diversity(D) = \min_{q \in Q} \left(\frac{1}{n} \sum_{i=1}^n f(q, D_i) \right)$$

where D is the anonymised dataset, Q is the set of all possible values for the sensitive attribute, n is the number of equivalence classes, D_i is the i -th equivalence class, and $f(q, D_i)$ is the proportion of records in the equivalence class D_i with a sensitive attribute value of q .

t-closeness t-closeness[16] aims to protect against attribute disclosure attacks by minimising the likelihood of inferring sensitive information based on background knowledge. In t-closeness, the notion of closeness refers to the similarity between the distribution of sensitive attributes in the original data and their distribution in the published data. The parameter "t" represents a threshold value that determines the acceptable level of similarity. A smaller value of t indicates a higher degree of privacy protection. t-closeness can be measured by Kullback-Leibler[10] Divergence. Kullback-Leibler divergence, also known as relative entropy, is a measure used to quantify the difference between two probability distributions. KL divergence measures the information loss when using one probability distribution Q to approximate another distribution P, given that P is the true distribution. Specifically, for two probability distributions P and Q, the KL divergence is defined as follows[7]:

$$KL(P \parallel Q) = \sum P(x) \log \left(\frac{P(x)}{Q(x)} \right)$$

Where $P(x)$ and $Q(x)$ represent the probabilities of event x under probability distributions P and Q, respectively.

To enhance privacy paradigms, common data anonymization operations include[5]:

- **Generalisation**[26]: In this operation, the original values of quasi-identifiers are transformed into less specific but semantically consistent values. For example, age or income can be generalised into intervals.
- **Suppression**[26]: This operation hides the original values of quasi-identifiers by replacing them with a special value. For instance, the value "20" of an individual's age can be anonymised. The value "0" can be replaced with "*", resulting in "2*" as the suppressed value of the quasi-identifier.
- **Perturbation**[26]: In this operation, random noise is added to the data to obscure the true values of individuals. This can be achieved by adding random numbers to numerical data or introducing randomisation processes in categorical data.
- **Microaggregation**[3]: This operation involves aggregating the data by combining multiple individuals' data into representative values, thus concealing specific individual information.

These anonymization techniques aim to balance the privacy protection and data utility in the published dataset. By applying these operations, data publishers can enhance the privacy of the dataset while preserving its usefulness for analysis and research purposes.

2.3 Federated Learning

Federated Learning[13][22] is a distributed machine learning approach that aims to train models without the need to send data from local devices to a central server. In traditional machine learning approaches, data is typically centralised at a single location for training, which involves data transmission and storage, posing security and privacy risks. The main idea of federated learning is to move the training process of models to local devices such as smartphones, tablets, or IoT devices. Each device locally stores its data and performs the model training process locally. Only the updated model parameters are sent to a central server for aggregation to update the global model. This approach keeps individual data locally without the need to share it with third parties or store it centrally, thereby enhancing data privacy and security. Federated learning offers several advantages. It can handle distributed, sensitive, and large-scale datasets while preserving user privacy. It is applicable in various scenarios such as healthcare, IoT, mobile devices, and edge computing. Through federated learning, individual devices can contribute their data to improve the performance of the global model while maintaining data confidentiality and privacy[11]. However, federated learning also poses several key challenges, such as inferred attacks on gradients, expensive communication costs between servers and clients, and device variability[17][6][32].

3 Methodology and framework

This section first describes our proposed framework for privacy-preserving federated learning, followed by an description of dataset and methods for anonymisation. Our approach aims to improve the anonymity of data and to evaluate the

impact of unionisation on the performance of federated learning and centralised machine learning models.

3.1 Data Description

In our study, we want to evaluate the impact of our framework on the performance of anonymity and federated learning models from a real, distributed dataset. Thus we utilise the Aposemat IoT-23 dataset[24] which aims to provide researchers with a large-scale, labelled dataset of IoT traffic to facilitate the development of machine learning algorithms. With 23 sub-datasets, the IoT-23 dataset covers a wide range of scenarios where network data was collected. These sub-datasets consist of network traffic data in pcap format, accompanied by labels indicating instances of malicious behaviour.

The dataset consists of 23 features, among which 'proto','orig_p' and 'ts' are selected as quasi-identifiers, and 'detailed-label' is sensitive data. 'proto' represents the protocol used in the network traffic packets, including 'tcp', 'udp', and 'icmp'. 'orig_p' represents the port used in the network traffic and 'ts' denotes the timestamp of the network connections, which has been standardised and retained with 4 significant digits. 'detailed-label' refers to the type of attack the system is subjected to.

3.2 Anonymisation methods

In order to enhance data anonymity, we employ two strategies: generalisation and microaggregation.

Generalisation Starting from the highest level of generalisation for the quasi-identifiers, we recursively specialise the partitions using multi-dimensional cuts until no further cuts can be made. In each iteration of the algorithm, a dimension (attribute) is selected for cutting. A common approach is to choose the dimension with the widest value range. Then, using the median split, we determine the splitting value and perform the cut based on that value.

Microaggregation Initially, we create clusters with at least k similar records. We then choose a representative value to replace all the quasi-identifiers within each group. The selection of the representative value can be done using different methods, such as the mean, median, or random selection.

3.3 A dual-layer privacy-preserving federated learning framework

Given that anonymization and federated learning achieve privacy protection at different stages of data processing, we propose a dual-layer framework (shown in Figure 1) that combines these two techniques to achieve comprehensive privacy protection from data collection to data utilisation.

The framework divides data processing into two stages: privacy data collection and privacy data usage.

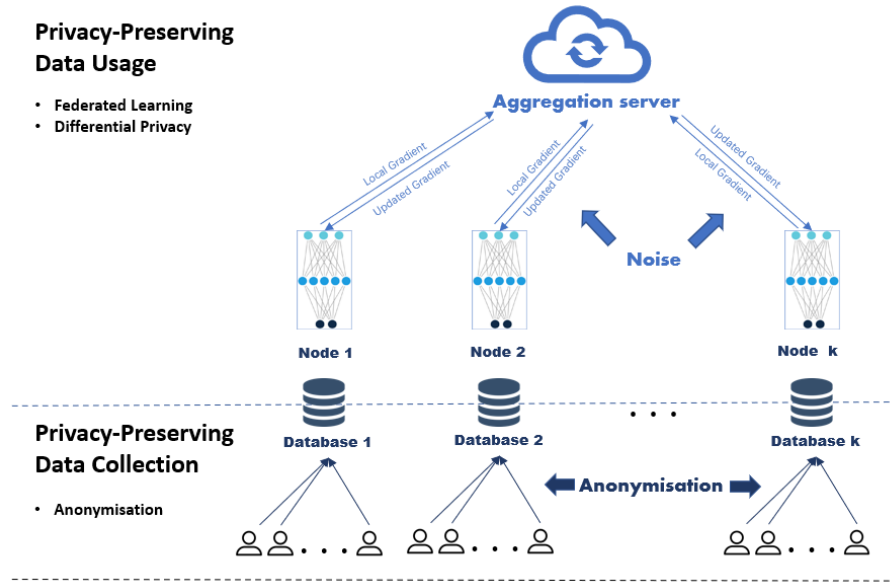


Fig. 1: Framework of the Dual-Layer Privacy-Preserving Federated Learning Framework.

1. In the first stage, after collecting individual information, the framework assesses the privacy metrics of the raw data and performs anonymization operations, such as generalisation and microaggregation, to ensure that data collection meets the requirements of Privacy-Preserving Data Publishing (PPDP).
2. In the second stage, through federated learning, the model is executed on multiple local datasets stored on various local nodes, utilising a central server to coordinate the training process. Additionally, the parameter update and aggregation processes between local nodes and the central coordinating server can be enhanced with differential privacy to strengthen data privacy protection. We add artificial noise to the parameters on the clients before aggregation to prevent inference attacks against the model gradients.

3.4 Privacy against centralisation

We generated two types of datasets, namely multiple distributed datasets and a single centralised dataset. First we sample 23 datasets from each of the 23 nodes as datasets $D\{a_1, a_2, \dots, a_{23}\}$ for distributed scenario; then we aggregate all the datasets from \mathbf{D} into a single dataset \mathbf{C} as the centralised scenario.

To explore the performance of data anonymization and federated learning, we consider the following scenarios:

Algorithm 1 Dual-Layer Privacy-Preserving Federated Learning

Stage 1: Privacy Data Collection**Input** : Original datasets $D\{d_1, d_2, \dots\}$.**Output**: Anatomised datasets $D'\{d'_1, d'_2, \dots\}$.**for** each dataset d in D **do** Identify the quasi-identifiers in d Measuring privacy metrics k -anonymity**for** each quasi-identifierApply generalisation or suppression to achieve k -anonymity**Return** the anatomised data set D' **Stage 2: Privacy Data Usage****Input** : Anatomised datasets D' , Mini-batch size (B),Participants per epoch (m), Total epochs (E), noise multiplier n .**Output**: Global model W_{GM} .**Aggregation Service Execution:****Initialise** W_{GM} :**for** each epoch =1,2,3 ... E **do** $D'_t \leftarrow$ (random set of m clients from C)**for** each participant d' in D' **do** $w_{GM}d'^{t+1} \leftarrow$ Update ($d', w_{GM}d'^t, n$) $w_{t+1} \leftarrow \sum_1^{d'} \frac{m_{d'}}{m} w_{GM}d'^{t+1}$ (Averaging Aggregation)**Client Update:** $\beta \leftarrow$ mini-batches creates through splitting local datasets D_L **for** each epoch =1,2,3 ... E **do****for** local mini-batch $b \in \beta$ **do** $w_{GM} \leftarrow w - \eta \Delta l(w, b)$ (Δl is the gradient of l on b and η is the learning rate)

-
1. Calculating different privacy parameters (k-anonymity, l-diversity, t-closeness) for a real-world dataset, IoT-23, in both centralised and federated learning distributed settings.
 2. Comparing the effects of different privacy-preserving anonymization strategies on federated learning and centralised machine learning models. To test the impact of data anonymization on the performance of federated learning and centralised machine learning models, firstly, we established three experimental scenarios:
 - (a) Federated Learning scenario: Anonymisation of the distributed datasets \mathbf{D} followed by federated learning.
 - (b) Centralised scenario: Anonymisation of the centralised dataset \mathbf{C} followed by centralised machine learning.
 - (c) Hybrid scenario: Anonymisation of the distributed datasets \mathbf{D} followed by aggregation of the datasets into a single consolidated dataset for centralised machine learning.

And then different levels of anonymization were implemented in each scenarios using generalisation and aggregation. Specifically, we tested anonymization with different degrees of k-anonymity (k-anonymity = 5, 10, 50, 100, 150, 200).

- Exploring the trade-off between data privacy and model performance in the proposed dual-layer privacy-preserving federated learning framework. We first selected the anonymization technique that resulted in the least performance loss for the federated learning model from the previous experiment. Then, we tested the performance of the models by varying the levels of anonymization in the data collection and the levels of noise in the federated learning to find a balance between data privacy and model performance.

4 Experiments

4.1 Privacy metrics

First, we measured the privacy metrics (K-anonymity, l-diversity, t-closeness) of the federated learning dataset and the centralised machine learning dataset on IoT-23. A higher value of K-anonymity and l-diversity and a lower value of t-closeness indicate a higher level of anonymization and better privacy of the data. Table 2 presents the results of the privacy metrics, and we can observe that the federated learning distributed dataset has lower values for each privacy metric compared to the centralised machine learning dataset.

	distributed dataset	Centralised dataset
k-anonymity	4	93
l-diversity	2	7
t-closeness	5.9	0.34

Table 2: Anonymity metrics of distributed datasets and centralised dataset

4.2 The impact of anonymization on model performance

We then compare the impact of different privacy-preserving anonymisation strategies on federated and centralised machine learning models, using generalisation, microaggregation, and increasing the k-anonymity of the data to 5, 10, 50, 100, 150 and 200. Figure 2 shows the accuracy of the models for each combination, where each bar chart represents a type of anonymization operation. The blue colour represents the accuracy of federated learning, the orange colour represents the accuracy of the hybrid scenario, and the green colour represents the accuracy of the centralised scenario.

The results show that under the anonymity constraints we tested, regardless of the anonymization method used, the performance of centralised machine learning only slightly decreases from an initial accuracy of 92% to a minimum of 89%. For federated learning and the hybrid scenario, as the anonymity constraints increase, the model’s accuracy continuously decreases, with federated

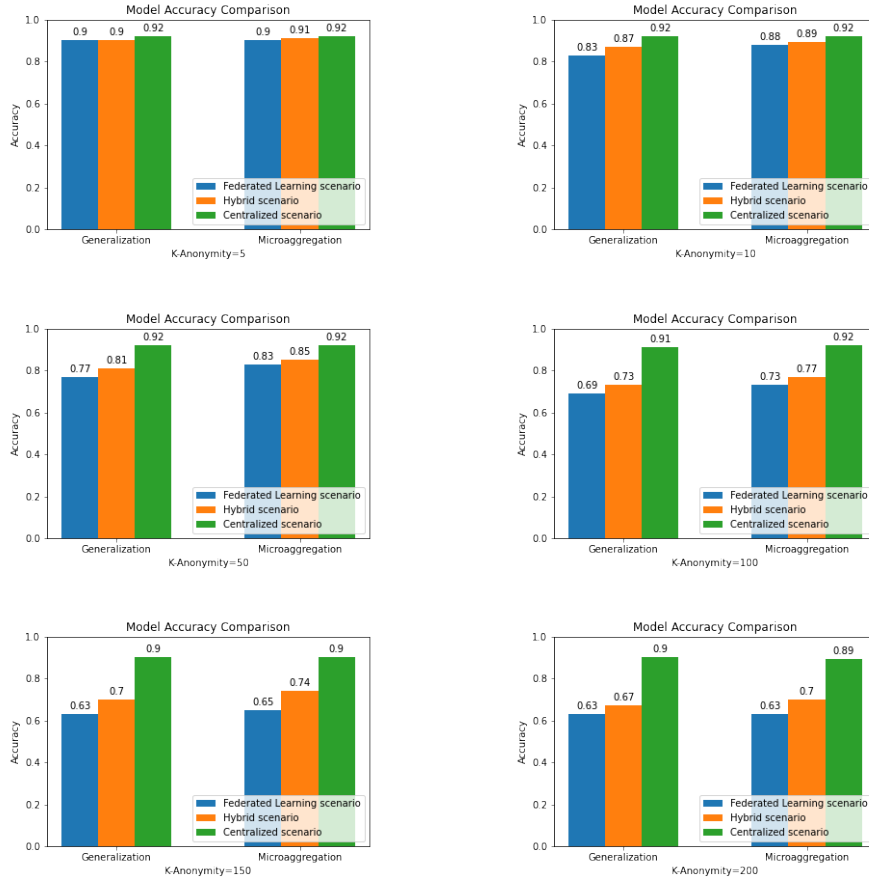


Fig. 2: Accuracy of the model after anonymisation

learning experiencing a larger decrease. Specifically, when the k-anonymity is less than or equal to 50, the Microaggregation operation incurs the least performance loss, with accuracy of 83% for federated learning and 85% for the hybrid scenario, which are still within an acceptable range. However, when the k-anonymity is greater than 100, the accuracy of federated learning drops significantly, reaching a minimum of 63%.

4.3 Trade-off in the Dual-Layer Framework

Based on the results of the previous experiment, in the data collection phase, we applied the Microaggregation strategy to anonymise the data to different levels (k-anonymity = 5, 10, 15). In the federated learning phase, we introduced

differential privacy by adding Gaussian noise with different noise multipliers (0.5, 1.0, 1.5) during gradient computation.

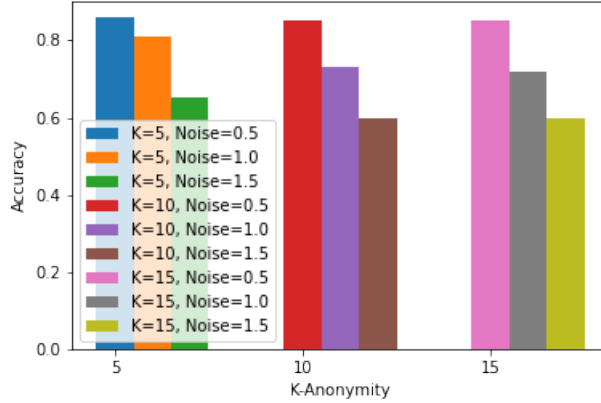


Fig. 3: Accuracy of different privacy strategies combinations.

Figure 3 shows the accuracy of federated learning under different privacy strategies. The results indicate that when the minimum anonymity constraint for data collection is set to 5 and the minimum noise factor is set to 0.5, the model achieves the highest accuracy of 85%. On the other hand, when using the highest level of privacy protection with a k-anonymity of 15 and a noise factor of 1.5, the model achieves the lowest accuracy of 61%. By observing the changes in accuracy, we can see that when the anonymity constraint for data collection is low, increasing data anonymity has less impact on the model’s accuracy. However, during the model training process, the noise factor of differential privacy has a significant impact on the model’s accuracy.

4.4 Evaluation and discussion

In our work, we compared the privacy metrics (k-anonymity, l-diversity, t-closeness) between centralised machine learning datasets and distributed federated learning datasets. The results showed that the anonymity of centralised machine learning datasets was better than that of federated learning datasets. This is because in federated learning, each node only stores local data, resulting in a smaller dataset with fewer quasi-identifiers, thus leading to lower anonymity.

Additionally, we applied different anonymization techniques to federated learning, centralised machine learning, and hybrid scenarios. We found that anonymization operations had a minor impact on centralised machine learning models, while significantly affecting the accuracy of federated learning models and hybrid scenario, with the greatest impact observed on federated learning models.

This is because the initial anonymity of the distributed data is lower, and achieving the same level of anonymity requires a more significant transformation of the distributed dataset, which further complicates the training of federated learning models. Among the tested anonymization techniques, microaggregation resulted in the least loss of model accuracy.

Finally, we combined privacy protection in data collection with differential privacy support in federated learning, resulting in a new dual-layer privacy protection framework. The results confirmed that within a certain range of anonymization operations (k -anonymity < 15) and the introduction of small noise perturbation (noise factor < 1), the framework could achieve dual-layer protection for data collection and utilisation while still maintaining an acceptable level of accuracy.

5 Conclusions and future work

This paper first compares the anonymity matrices between federated learning distributed datasets and centralised machine learning datasets using real-world datasets. It is found that the federated learning datasets have lower anonymity compared to the centralised machine learning dataset. Furthermore, the impact of different privacy-preserving anonymization strategies on federated learning and centralised machine learning models is evaluated. The results indicate that each anonymization operation leads to a decrease in the accuracy of the federated learning models compared to the centralised machine learning models. When the anonymity constraints are relatively large, the model accuracy can drop to below 70%. This is because the initial anonymity of the federated learning data is low, and a greater degree of data transformation is required to achieve the same level of anonymity. Among the tested anonymization techniques, Microaggregation exhibits the least loss in model accuracy.

Moreover, a dual-layer privacy-preserving federated learning framework is proposed. In the data collection phase, the Microaggregation strategy is applied for anonymising the data, while in the federated learning phase, differential privacy is achieved by adding Gaussian noise with different noise multipliers during the gradient computation. Evaluation on real-world datasets demonstrates the improvement in data anonymity and model performance achieved by this framework. Through these contributions, a comprehensive understanding of the impact of anonymization strategies on federated learning is gained, filling a gap in the field of privacy protection.

In the future, we will explore the impact of other anonymization strategies on federated learning, such as Top-Down Greedy Anonymisation and k -NN Clustering-Based Anonymisation. Additionally, we will compare various privacy metrics to select and apply them based on specific privacy protection requirements and scenarios, aiming to optimise the effectiveness of privacy protection measures.

References

1. Ayala-Rivera, V., McDonagh, P., Cerqueus, T., Murphy, L., et al.: A systematic comparison and evaluation of k-anonymization algorithms for practitioners. *Transactions on data privacy* **7**(3), 337–370 (2014)
2. Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., Kiddon, C., Konečný, J., Mazzocchi, S., McMahan, B., et al.: Towards federated learning at scale: System design. *Proceedings of machine learning and systems* **1**, 374–388 (2019)
3. Domingo-Ferrer, J., Mateo-Sanz, J.M.: Practical data-oriented microaggregation for statistical disclosure control. *IEEE Transactions on Knowledge and data Engineering* **14**(1), 189–201 (2002)
4. Fung, B.C., Wang, K., Chen, R., Yu, P.S.: Privacy-preserving data publishing: A survey of recent developments. *ACM Computing Surveys (Csur)* **42**(4), 1–53 (2010)
5. Gkoulalas-Divanis, A., Loukides, G., Sun, J.: Publishing data from electronic health records while preserving privacy: A survey of algorithms. *Journal of biomedical informatics* **50**, 4–19 (2014)
6. Hao, M., Li, H., Xu, G., Liu, S., Yang, H.: Towards efficient and privacy-preserving federated deep learning. In: *ICC 2019-2019 IEEE international conference on communications (ICC)*. pp. 1–6. IEEE (2019)
7. Hershey, J.R., Olsen, P.A.: Approximating the kullback leibler divergence between gaussian mixture models. In: *2007 IEEE International Conference on Acoustics, Speech and Signal Processing-ICASSP'07*. vol. 4, pp. IV–317. IEEE (2007)
8. Jayabalan, M., Rana, M.E.: Anonymizing healthcare records: a study of privacy preserving data publishing techniques. *Advanced Science Letters* **24**(3), 1694–1697 (2018)
9. Ji, Z., Lipton, Z.C., Elkan, C.: Differential privacy and machine learning: a survey and review. *arXiv preprint arXiv:1412.7584* (2014)
10. Joyce, J.M.: Kullback-leibler divergence. In: *International encyclopedia of statistical science*, pp. 720–722. Springer (2011)
11. Kairouz, P., McMahan, H.B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A.N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., et al.: Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning* **14**(1–2), 1–210 (2021)
12. Kanwal, T., Anjum, A., Khan, A.: Privacy preservation in e-health cloud: taxonomy, privacy requirements, feasibility analysis, and opportunities. *Cluster Computing* **24**, 293–317 (2021)
13. Konečný, J., McMahan, H.B., Yu, F.X., Richtárik, P., Suresh, A.T., Bacon, D.: *Federated learning: Strategies for improving communication efficiency* (2017)
14. Langheinrich, M.: Privacy by design—principles of privacy-aware ubiquitous systems. In: *UbiComp 2001: Ubiquitous Computing: International Conference Atlanta Georgia, USA, September 30–October 2, 2001 Proceedings*. pp. 273–291. Springer (2001)
15. Li, L., Fan, Y., Tse, M., Lin, K.Y.: A review of applications in federated learning. *Computers & Industrial Engineering* **149**, 106854 (2020)
16. Li, N., Li, T., Venkatasubramanian, S.: t-closeness: Privacy beyond k-anonymity and l-diversity. In: *2007 IEEE 23rd international conference on data engineering*. pp. 106–115. IEEE (2006)

17. Li, T., Sahu, A.K., Talwalkar, A., Smith, V.: Federated learning: Challenges, methods, and future directions. *IEEE signal processing magazine* **37**(3), 50–60 (2020)
18. Liu, B., Ding, M., Shaham, S., Rahayu, W., Farokhi, F., Lin, Z.: When machine learning meets privacy: A survey and outlook. *ACM Computing Surveys (CSUR)* **54**(2), 1–36 (2021)
19. Machanavajjhala, A., Kifer, D., Gehrke, J., Venkatasubramanian, M.: l-diversity: Privacy beyond k-anonymity. *ACM Transactions on Knowledge Discovery from Data (TKDD)* **1**(1), 3–es (2007)
20. Majeed, A., Lee, S.: Anonymization techniques for privacy preserving data publishing: A comprehensive survey. *IEEE access* **9**, 8512–8545 (2020)
21. Majeed, A., Lee, S.: Anonymization techniques for privacy preserving data publishing: A comprehensive survey. *IEEE Access* **9**, 8512–8545 (2021). <https://doi.org/10.1109/ACCESS.2020.3045700>
22. McMahan, B., Moore, E., Ramage, D., Hampson, S., y Arcas, B.A.: Communication-efficient learning of deep networks from decentralized data. In: *Artificial intelligence and statistics*. pp. 1273–1282. PMLR (2017)
23. Nergiz, M.E., Atzori, M., Clifton, C.: Hiding the presence of individuals from shared databases. In: *Proceedings of the 2007 ACM SIGMOD international conference on Management of data*. pp. 665–676 (2007)
24. Parmisano, A., Garcia, S., Erquiaga, M.J.: A labeled dataset with malicious and benign iot network traffic. *Stratosphere Laboratory: Praha, Czech Republic* (2020)
25. Regulation, P.: Regulation (eu) 2016/679 of the european parliament and of the council. *Regulation (eu)* **679**, 2016 (2016)
26. Samarati, P.: Protecting respondents identities in microdata release. *IEEE transactions on Knowledge and Data Engineering* **13**(6), 1010–1027 (2001)
27. Slijepčević, D., Henzl, M., Klausner, L.D., Dam, T., Kieseberg, P., Zeppelzauer, M.: k-anonymity in practice: How generalisation and suppression affect machine learning classifiers. *Computers & Security* **111**, 102488 (2021)
28. Sweeney, L.: k-anonymity: A model for protecting privacy. *International journal of uncertainty, fuzziness and knowledge-based systems* **10**(05), 557–570 (2002)
29. Vedangi, A., Anandam, V.: Data slicing technique to privacy preserving and data publishing. *Cancer* **4790**(4790), 4790 (2013)
30. Wei, K., Li, J., Ding, M., Ma, C., Yang, H.H., Farokhi, F., Jin, S., Quek, T.Q., Poor, H.V.: Federated learning with differential privacy: Algorithms and performance analysis. *IEEE Transactions on Information Forensics and Security* **15**, 3454–3469 (2020)
31. Wieringa, J., Kannan, P., Ma, X., Reutterer, T., Risselada, H., Skiera, B.: Data analytics in a privacy-concerned world. *Journal of Business Research* **122**, 915–925 (2021)
32. Yang, H.H., Arafa, A., Quek, T.Q., Poor, H.V.: Age-based scheduling policy for federated learning in mobile edge networks. In: *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 8743–8747. IEEE (2020)