Future Technology
Open Access Journal
ISSN 2832-0379

May 2026| Volume 05 | Issue 02 | Pages 179-188

Journal homepage: https://fupubco.com/futech

https://doi.org/10.55670/fpll.futech.5.2.17

Future Publishing LLC

Article

Federated learning for breaking data silos in smart governance: a privacy-preserving framework for cross-agency collaboration

Que Zhang*

School of Cross-border E-commerce, Yango University, Fuzhou 350015, China

ARTICLE INFO

Article history:
Received 29 October 2025
Received in revised form
28 December 2025
Accepted 03 February 2026

Keywords:
Federated learning, Smart governance,
Privacy protection, Differential privacy, Data silos

*Corresponding author
Email address:
m15942983900@163.com

DOI: 10.55670/fpll.futech.5.2.17

ABSTRACT

The realization of smart governance highly relies on the effective integration and collaborative utilization of cross-departmental government data, yet data silos that have formed over time and the privacy compliance risks faced by traditional centralized sharing models severely constrain the improvement of collaboration effectiveness. Addressing this challenge, this study proposes the FedGov privacy-preserving federated learning framework for smart governance scenarios, designing a three-layer system architecture comprising data, computation, and coordination layers to support multi-departmental heterogeneous data collaboration, and developing the FedGov-DP algorithm integrating dual mechanisms of differential privacy and secure aggregation to realize the "data usable but invisible" cross-departmental collaboration paradigm. Systematic experiments simulating government scenarios based on public datasets demonstrate that the proposed framework effectively breaks down data silos and achieves significant collaboration gains, the differential privacy mechanism effectively defends against membership inference attacks, and the method exhibits good adaptability to moderate data heterogeneity common in government scenarios. This study extends the application boundaries of federated learning in the public governance domain, provides a new technical pathway for addressing the government data silo dilemma, and the constructed framework with parameter configuration guidelines offers technical support and practical reference for smart governance digital transformation.

1. Introduction

Smart governance, as the main goal of digital government construction, heavily depends on the efficient integration and joint use of cross-departmental government data for its realization. With the background of the in-depth development of government service reforms like "One-Stop online services" and "cross-provincial services," overcome the barriers of data between departments to realize precise public services, which has become an important approach to improve governance efficiency. The unclear positioning of role relationships and responsibility boundaries in cross-departmental collaboration has a heavy negative effect on the intensity and scale of data exchange [1]. The reason why the phenomenon of data silos actually arose in government affairs is not only because of the complex information systems developed independently by different departments across different periods of time, but is also closely linked to the risks of privacy compliance for the conventional centralized data sharing model. Especially for such sensitive areas as taxation, social security, and medical affairs, direct acquisition of original data can give rise to problems of leakage of privacy among citizens as well as disputes regarding data sovereignty. The improvement of data sharing for digital government needs to find a balance between openness and security [2]. There is a paradigm shift in the governance approach for interagency collaboration from bureaucratic to a network-centric approach in the collaborative governance paradigm. This results in more intensity of data fusion required from the underlying technological platform [3]. There are two tasks in the design of a secure framework for open government data, including both its availability and the protection of privacy [4]. As a new pattern of distributed machine learning of "data stays stationary while the model moves," FL offers a new technical solution for the partnership dilemma of government data, where the training of the model can be done locally without sharing data. On the security level, the analysis of federated learning threats has established a systematic theory with regard to data poisoning attacks, model poisoning attacks, and inference attacks, which has provided a theoretical basis for designing defenses [5]. On the

system level, the study of technical challenges with regard to communication efficiency optimization, heterogeneous processing, and scalability has moved federated learning forward from theory to practice [6]. The target of algorithm optimization has been expanded from traditional FedAvg to other cutting-edge areas like data heterogeneity adaptation and privacy improvement [7]. Privacy and security are the most important issues of federated learning. The integration of technologies on differential privacy and secure multi-party computation has been fully discussed [8], while the balance between model utility and budget in privacy mechanisms has also been the key point of research [9]. From the implementation level, the corresponding security guarantee mechanism for federated learning systems is required to be a full chain, involving the encryption of communication, access control, and traceability for audit reasons [10], and the research view is increasingly focusing on the dimension of interpretability and fairness within trusted artificial intelligence [11]. From the application level, the federated learning has proved its dual roles of either preserving the privacy or collectively modeling in the sensor network of a smart city [12], and its comprehensive application with the Internet of Things is increasingly being enhanced [13]. However, there are three gaps in existing research on the application of intelligent governance in a systematic way. Firstly, existing research is mainly based on business cases like healthcare and finance, while it rarely takes into consideration specifics like the relationship of rights and responsibilities between departments, as well as differences of sensitivity levels of data on cross-departmental collaboration of government data. Secondly, existing research on protection of privacy is mainly based on a single technical approach, without a system architecture combining differential privacy and security aggregation. Thirdly, existing research on algorithm adjustment and verification on Non-IID properties of data related to cases of government affairs is less adequate. Such constraints, taken together, point to the urgent demand for a holistic approach capable of addressing the unique institutional and technological imperatives associated with privacy-preserving collaboration on data in the context of public governance.

To address the above research gaps, in this study we aspire to build a privacy-preserving federated learning framework for smart governance, focusing on joint modeling of inter-departmental government data through the combination of differential privacy and secure aggregation. The research value of this work is embodied in three ways. Firstly, from an architectural perspective, a federated learning framework suitable for multi-departmental heterogeneous data collaboration in government affairs is proposed. Secondly, from an algorithmic perspective, a federal optimization algorithm integrating multiple privacy protection measures is put forward in order to offer measurable privacy protection. Thirdly, from an empirical perspective, simulation studies of inter-departmental collaboration scenarios are conducted based on real-world datasets, and the validity of the framework is demonstrated from three aspects of collaboration efficiency, privacy protection, and scenario adaptability. Theoretical research value of this research is that it breaks through the application limits of federated learning in public administration. Practical research value of this research is that it provides a "usable but invisible" technology solution for inter-departmental data sharing of governments. It seeks to unleash the efficiency of government data in governance.

2. Methodology

2.1 System architecture design

The underlying aim of smart governance is to reach precision in public service provision via collaborative data across various departments, while data silos established over time between departments impede it greatly. To model the challenge of collaborative data across departments, consider K government departments serving as participants in the federated learning system, where the k-th department $(k = 1, 2, \ldots, K)$ holds a local dataset D_k containing n_k samples. The collaboration objective is to train a globally optimal model w^* without raw data leaving departmental boundaries, which can be formally expressed as minimizing the weighted empirical loss function:

$$w^* = arg\min_{w} F(w) = arg\min_{w} \sum_{k=1}^{K} \frac{n_k}{n} F_k(w) \quad (1)$$

$$F_k(w) = \frac{1}{n_k} \sum_{j=1}^{n_k} \ell(w; x_j^{(k)}, y_j^{(k)}) \quad (2)$$

where $F_k(w)$ represents the local empirical loss of department k, $\ell(\cdot)$ denotes the sample-wise loss function, $(x_j^{(k)}, y_j^{(k)})$ denotes the j-th data sample in department k, and $n = \sum_{k=1}^{K} n_k$ represents the total sample size across all departments. It is evident in this optimization problem that the value of the data brought by each department is critical in optimizing the performance of the model, while at the same time being mindful of the fact that sensitive information such as tax and social security records must remain within the boundaries of the respective departments. In this problem, the data value provided by the departments is essential for improving the model, and at the same time, the problem ensures no leakage of private information like taxation information and social security information. Data aggregation approaches using a centralized data model face implementation difficulties due to sovereignty concerns and privacy compliance [14], urgently demanding the exploration of technical routes for "data usable but invisible" to overcome the data silo problem.

To meet these issues, this paper presents the FedGov three-layer system architecture. The data tier will be conducted in each of the involved departments, performing in-house operations regarding storage, value filling, encoding, and other preprocessing tasks with the focus that no original citizen data will be transferred beyond the boundaries of each administration, thus addressing issues of sovereignty concerning data aggregation on the institutional level. However, it is important to incorporate computation on model parameters by considering mini-batch stochastic gradient descent on the computation level, which is implemented on the same level as the data level within departments, utilizing methods like gradient clipping and adding noise to prevent reverse-engineering of transmitted data beyond department boundaries. The coordination level is implemented on the governmental cloud system, performing aggregated gradient computation on encrypted model updates provided by each department, producing global model parameters while sending copies back to each department to begin a new round of computation. This level of coordination is incapable of accessing either raw or unprotected model update data, acting purely as a middleman or coordinator facilitating collaboration among departments while promoting cooperation as opposed to relying on trust

among departments. To provide an intuitive understanding about the organizational relationship and information flow for the three-layer architecture, the design for the FedGov system is given below in Figure 1.

Figure 1 shows how this architecture enables a clean division of labor through hierarchical design, where the data layer preserves data sovereignty across departments by having a computation layer perform localized model training and privacy protection, with a coordination layer carrying out secure aggregation to produce a global model that never sees the original data. This hierarchical architecture, as discussed above, must also handle the heterogeneous nature of the data for government affairs. Every department is different based on functionalities, and the service scope also varies, resulting in imbalanced label distributions and attribute data collected based on the functionalities of the departments, representing the Non-IID setting. The Dirichlet Distribution-based partitioning strategy with the value of α is also taken as the standard for simulating the Non-IID setting for federated learning studies. The value of α determines the level of label distribution heterogeneity among departments, with low values representing high heterogeneity in which a small number of classes dominate each department's data, and high values representing a move towards IID settings with a relatively even distribution of labels. The Non-IID factors introduced can result in local models converging to department-wise optimal points [15], and are tackled in the design of the algorithm in the next section.

2.2 Privacy-preserving federated learning algorithm

The Federated Averaging (FedAvg) algorithm set a precedent for Privacy Preserving Distributed Learning under the "data stays local, model moves" framework [16], where each party locally trains on local data and sends only model updates while the global model is computed by weighted averaging by the server. However, FedAvg is known to lack formal privacy protection and is susceptible to attacks like membership inference attacks on model parameters [17]. On the basis of this foundation work, the paper presents the FedGov-DP algorithm that meets the key sensitivity requirements for publicly shared data based on both concepts of differential privacy and secure aggregation.

Unlike the DP-FedAvg algorithm, which supports the concept of differential privacy alone, the FedGov-DP algorithm also supports the concept of secure aggregation that is required for the specific trust issues associated with the cooperation of government departments. The FedGov system assumes the honest-but-curious adversary model, which is a commonly adopted model for privacy-preserving distributed systems. This model captures the realistic setting of government agency collaborations, where the departments are institutionally honest and follow the protocols, but might have reasons to inquire about the data distributions of other departments. Under the above threat model, the departments and the coordination level are expected to follow the protocol as defined by the algorithm, without any divergence from the algorithm specification, while they might try to derive private information from the received messages. The coordination level is able to see the incoming model updates, but it cannot see the raw gradients based on the secure aggregation process that will be discussed below, and the departments cannot see the data and the unmasked model updates of the other departments. There could be membership inference attacks from external attackers not belonging to the federation based on auxiliary datasets that have the same distributions, and the framework does not include Byzantine attackers that could send malicious updates arbitrarily, which is recognized as a limitation for future studies.

The algorithm iterates through communication rounds $t=1,\ldots,T$: the coordination layer distributes the global model w^t, each department uploads privacy-protected updates after completing local training, and the coordination layer aggregates to generate w^{t+1}. During the local training phase, department k executes E rounds of mini-batch gradient descent on private data. Local multi-step updates can significantly reduce communication complexity [18], but may cause client drift in Non-IID scenarios. This algorithm constrains update magnitude through subsequent privacy mechanisms to balance efficiency and stability. After training completion, the update is computed as $\Delta w_k^t = w_k^{t,E} - w^t$.

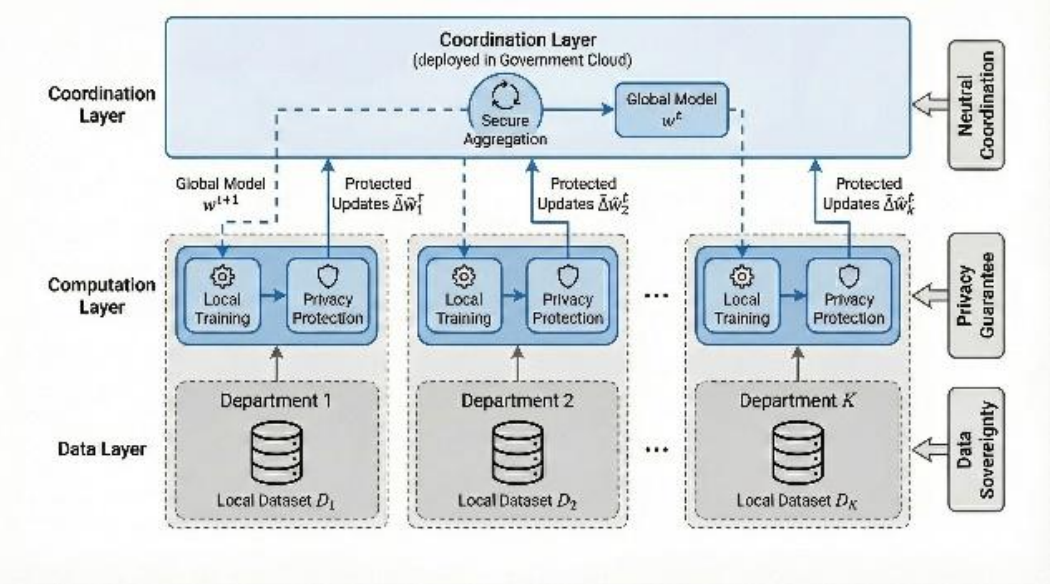

Figure 1. FedGov Three-Layer System Architecture

The differential privacy mechanism ensures the privacy of the model updates. (ε, δ)-differential privacy guarantees the bound on the difference of the output distribution due to the addition or deletion of a single example [19], where lower ε represents a better privacy level. Privacy protection uses a two-step gradient perturbation approach. Gradient clipping limits the L_2 norm of the updates below a threshold C:

$$\overline{\Delta w_k}^t = \Delta w_k^t / max\left(1, \frac{||\Delta w_k^t||_2}{C}\right) \quad (1)$$

This step determines the upper bound of the sensitivity of the mechanism. It is a mandatory mathematical step leading to the noise calibration. Adding noise injects the calibrated Gaussian noise:

$$\Delta w_k^t = \overline{\Delta w_k}^t + \mathrm{N}\,(0, \sigma^2 C^2 I) \quad (4)$$

The standard deviation of the noise σ is determined jointly with the pair (ε, δ) and the sampling rate [20]. The Moments Accountant algorithm is employed because it offers better bounds in composing privacy budgets than other accounting techniques. The standard composition bounds are loose and linear or sub-linear in the number of calls to a mechanism, making them impractical in iterative learning settings because privacy guarantees become extremely loose with every iteration and are not useful in providing privacy protection in training models with hundreds of communication rounds. The Moments Accountant follows the moment generating function of privacy loss random variables in every iteration and thus facilitates better (ε,δ)-differential privacy bounds.

The secure aggregation algorithm prevents the cooperation layer from receiving anything except aggregated results, without having to view the contributions from each department. In the pairwise masking approach, each pair of departments establishes a shared random seed through key exchange, from which both departments generate identical random masks [21]. For each department k, the masked update is computed as $\widehat{\Delta w}_k^t \leftarrow \widetilde{\Delta w}_k^t + \sum_{j<k} r_{kj} - \sum_{j>k} r_{jk}$, where r_{kj} denotes the shared mask with department j. Since departments k and j share the same mask value but apply opposite signs, each mask term r_{kj} added by department k is canceled by the corresponding term- r_{jk} subtracted by department j when all updates are summed at the coordination layer. This guarantees that the coordinator obtains only the aggregate of original updates without observing individual contributions.

The secure aggregation protocol results in additional overhead compared to plaintext aggregation, but this overhead is only O(K²) key exchanges for K departments in the setup stage, which is amortized over all communication rounds. For every round of aggregation, the communication overhead for the uploaded masked updates is the same as that for the original model updates, as there is no additional overhead. The overhead for mask generation by pseudorandom number generators is negligible compared to local model training times, and for the setup used in this work with a limitation of K≤6 departments, this overhead is also acceptable.

Algorithm 1 is the complete pseudocode for the FedGov-DP solution, integrating differential privacy and secure aggregation to ensure end-to-end privacy for interdepartmental cooperation. The iterative process involves three main phases: local training with mini-batch stochastic gradient descent, differential privacy protection with gradient clipping and calibrated noise addition, and secure aggregation with pairwise masking. The algorithm ends after communication rounds, yielding the final global model that has learned from the knowledge of all interdepartmentally cooperating units while preserving the confidentiality of the data and the gradients throughout the cooperation process.

Algorithm 1: FedGov-DP

Input: Departments K, local datasets $\{D_k\}_{k=1}^K$, rounds T, local epochs E, clipping threshold C, noise scale σ

Output: Global model w^T

Initialize global model w^0

for round $t = 0$ to $T-1$ **do**

  Broadcast w^t to all departments

  for each department $k \in \{1, \ldots, K\}$ **in parallel do**

    Local Training

    $w_k^{t,E} \leftarrow \mathrm{LocalSGD}(w^t, D_k, E)$

    $\Delta w_k^t \leftarrow w_k^{t,E} - w^t$

    Differential Privacy Protection

    $\overline{\Delta w_k}^t \leftarrow \Delta w_k^t / \max\left(1, \frac{\| \Delta w_k^t \|_2}{C}\right)$ ▷ Gradient clipping

    $\Delta w_k^t \leftarrow \overline{\Delta w_k}^t + \mathrm{N}\,(0, \sigma^2 C^2 I)$ ▷ Noise injection

    Secure Aggregation

    $\Delta w_k^t \leftarrow \Delta w_k^t + \sum_{j<k} r_{kj} - \sum_{j>k} r_{jk}$ ▷ Pairwise masking

    Upload Δw_k^t to coordination layer

  end for

  Global Aggregation

  $w^{t+1} \leftarrow w^t + \sum_{k=1}^K \frac{n_k}{n} \Delta w_k^t$

end for

return w^T

2.3 Experimental design

Experiments are conducted on two publicly available datasets from the UCI Machine Learning Repository based on principles that would facilitate interdepartmental collaboration in simulations that are fully publicly accessible, exempt from ethics review approval, and widely used in federated learning literature. The Adult Census dataset (https://archive.ics.uci.edu/dataset/2/adult) has 48,842 U.S. census records with 14 attributes. There is a class distribution of 24.1% high-income people and 75.9% low-income people. The binary classification problem is the prediction of earned income above $50,000. For the simulation scenario, the precise configuration involves joint resident precision services from 6 departments: taxation, social security, civil affairs, human resources, housing, and statistics. It involves a Non-IID environment where the partitioning of the data involves the occupation attribute. The Credit Card dataset (https://archive.ics.uci.edu/dataset/350/default+of+credit+card+clients) has 30,000 credit card client records. It has 23 attributes. There is a class distribution of 22.1% defaulters and 77.9% non-defaulters. The binary classification problem involves the probability of default. The precise simulation involves joint credit default probability supervision from 5 institutions. Data is partitioned according to level of education. Table 1 provides the specifics of the precise configurations.

Table 1. Experimental dataset configuration

Dataset	Samples	Features	Class Distribution	Departments	Partitioning Strategy
Adult Census	48,842	14	24.1% (>50K) / 75.9% (≤50K)	6	By occupation attribute
Credit Card	30,000	23	22.1% (default) / 77.9% (non-default)	5	By education level

Table 2. Hyperparameter configuration and selection rationale

Parameter	Symbol	Value	Selection Rationale
Communication rounds	T	100	Sufficient for convergence based on pilot experiments
Local epochs	E	5	Balances computation and communication; higher values may exacerbate client drift in Non-IID settings
Batch size	B	32	Standard choice providing stable gradient estimates
Learning rate	η	0.01	Determined via grid search over candidate set
Clipping threshold	C	1.0	Controls gradient sensitivity based on typical gradient norms
Noise multiplier	σ	1.1	Calibrated to achieve target ε using Moments Accountant
Privacy budget	ε	4.0	Moderate privacy-utility trade-off; sensitivity analysis over ε∈{1,2,4,8}
Dirichlet parameter	α	0.5	Simulates moderate Non-IID conditions; sensitivity analysis over α∈{0.1, 0.5,1.0,5.0}

All baseline models consist of 8 strategies: classical algorithms FedAvg and Local SGD, and algorithms designed for addressing system and statistical heterogeneity in federated learning: FedProx [22], addressing client drift in federated learning using SCAFFOLD, addressing heterogeneous local updates using FedNova [23], addressing privacy in federated learning using DP-FedAvg [24], and finally upper and lower bounds given by central training and local training. Other privacy-preserving methods that were not considered are those that use pure secure multi-party computation without differential privacy, since these methods are cryptographically secure during the computation phase, though they are not formally secured in terms of inference attacks on the published model. Also, methods that use homomorphic encryption were ruled out due to their high computational costs that are impractical for deep learning. Local SGD is different from Local Only since it involves model synchronization among departments at certain intervals, while the model is trained individually without any communication for the Local Only method, and the purpose of the Local SGD algorithm is to provide a communication-efficient federated learning baseline, while the Local Only algorithm sets the lower bound for the model's performance, which signifies the absence of cooperation.

Evaluation metrics are divided into three categories: performance of the model on accuracy, F1 score, and AUC-ROC; privacy protection on cumulative privacy budget ε_{total} and membership inference attack success rate to measure empirical privacy leakage [25]; and collaboration gain as the relative performance improvements of federated learning over local training. The testbed setup will utilize NVIDIA RTX 3090 GPU, PyTorch 1.12, and the Flower 1.5 federated learning framework. The model architecture used in the task will be a two-layer fully connected neural network with a hidden state of size 128 and a ReLU activation function. The reason for choosing such a model architecture for the task is the tabular structure of the two datasets with 14 and 23 features, which suggests the use of a dense model over a convolutional model suited for images or a recurrent model suited for sequential data, since the size of the two datasets is also not large enough to justify the use of a deep model which can easily lead to overfitting in binary classification problems. Table 2 summarizes the hyperparameter configurations adopted in the experiments along with the selection rationale. As shown in Table 2, the privacy-related parameters are calibrated to balance protection strength and model utility. The value of the clipping threshold and the noise scale parameter is a determining factor in the level of protection offered by differential privacy. Gradient clipping is a mechanism to constrain the L2 norm of individual updates to prevent any particular data point from contributing more than a certain value to the update, thereby setting a bound on the sensitivity of the differential privacy algorithm. The scaled Gaussian noise is then applied to the clipped contribution, and the level of the noise is set based on the privacy budget and the sensitivity value.

3. Result

3.1 Validation of Federated Collaboration Effectiveness

The true power of cross-departmental data collaboration relies on the integration of fragmented information to improve the performance of models. This section verifies whether the proposed model, FedGov-DP, can be utilized to achieve the same while maintaining privacy. The design of the experiment is based on these principles: Local Only (independent learning of each department) is the performance lower bound and is for determining gains due to collaboration; Centralized Learning (training of data in the central department) is the performance upper bound and will determine how close federated learning is to the optimum solution; and DP-FedAvg and other approaches are the performance baselines for checking the appropriateness of optimization of FedGov-DP. Unless otherwise noted, the experiments follow Non-IID data splitting with Dirichlet parameter $\alpha = 1.0$, with a privacy budget $\varepsilon = 4$, and repeated with a set configuration 5 times with mean and standard

deviation. Table 3 presents the classification results of the classification problem with 9 models on the Adult Census dataset. As is clear from Table 3, the values for FedGov-DP with ε = 4 provide an accuracy of 84.2%, F1 score of 0.683, and AUC ROC value of 0.891, which improve over the Local Only method by 10.7% relative to the baseline accuracy of 73.5%, while the difference with Centralized Learning is held within 2.6%. Compared with the baseline models of non-private federated learning, FedGov-DP slightly deteriorates compared with FedAvg (85.6%) and SCAFFOLD (85.9%), due to the noise induced by privacy protection, but it still surpasses DP-FedAvg (82.1%) by a margin of 2.1 percentage points, proving the optimization function.

Table 3. Performance comparison of different methods on adult census dataset

Method	Accuracy (%)	F1-Score	AUC-ROC
Centralized	86.8 ± 0.31*†	0.721 ± 0.014*†	0.913 ± 0.007*†
SCAFFOLD	85.9 ± 0.47*†	0.706 ± 0.019*†	0.902 ± 0.011*†
FedAvg	85.6 ± 0.52*†	0.701 ± 0.016*†	0.898 ± 0.009*†
FedNova	85.2 ± 0.44*†	0.694 ± 0.021*†	0.903 ± 0.013*†
FedProx	84.8 ± 0.39*	0.689 ± 0.018*	0.892 ± 0.010*
Local SGD	84.3 ± 0.61*	0.685 ± 0.023*	0.889 ± 0.016*
FedGov-DP (ours)	84.2 ± 0.43*	0.683 ± 0.017*	0.891 ± 0.012*
DP-FedAvg	82.1 ± 0.53*	0.651 ± 0.019*	0.867 ± 0.014*
Local Only	73.5 ± 1.18	0.547 ± 0.031	0.794 ± 0.022

Note: * Statistically significant improvement over Local Only ($p < 0.05$, paired t-test). † Statistically significant improvement over FedGov-DP ($p < 0.05$).

For the purpose of checking the robustness of the foregoing conclusions across different scenarios, Table 4 lists the experimental results on the Credit Card dataset, which has more aggravated circumstances of class imbalance (roughly 22.1% for the default samples). As illustrated in Table 4, for the Adult Census data, the accuracy of 81.5%, F1-score of 0.614, and AUC-ROC of 0.867 for FedGov-DP is obtained, which is an improvement of 7.3% over Local Only (74.2%), while closing the gap of 1.6% to Centralized (83.1%), and performing better than DP-FedAvg (79.3%) by 2.2%. The results confirm similar trends and the applicability of the new approach to data of different characteristics and complexities. Convergence during the training procedure will directly influence the communication overhead and cost of time during actual deployment. To assess the efficiency of the training procedure of FedGov-DP, the convergence curve of dominant techniques on both datasets is shown in Figure 2, where the x-coordinate represents the number of communication rounds, and the y-coordinate represents the test set accuracy.

Figure 2(a) illustrates that FedGov-DP converges between the non-private baselines (FedAvg and SCAFFOLD) and DP-FedAvg on the Adult Census dataset with competitive training efficiency while ensuring data privacy. Figure 2(b) highlights that there are steady levels of convergence on the Credit Card dataset with less final accuracy because of class imbalance, which proves that the dual privacy approach ensures consistent model performance regardless of the dataset.

3.2 Validation of privacy protection effectiveness

The effectiveness validation on federated collaboration proves the efficiency improvement and privacy maintenance of FedGov-DP. In the following, the effectiveness on privacy protection is discussed based on two aspects: privacy utility trade-off and defense against attacks. Differential privacy manages the level of protection based on the privacy budget $\mathcal{E}$. The smaller the $\mathcal{E}$, the better the level of protection but the more noise, a trade-off for privacy and utility. Experiments on the two datasets respectively adopt four different parameters for $\varepsilon \in \{1,2,4,8\}$, among which $\varepsilon = 1$ is for highly sensitive tax and social security information, and $\varepsilon = 8$ is for statistical information. The results on the two datasets are presented in Figure 3.

Table 4. Performance comparison of different methods on credit card dataset

Method	Accuracy (%)	F1-Score	AUC-ROC
Centralized	83.1 ± 0.29	0.641 ± 0.019	0.881 ± 0.008
SCAFFOLD	82.4 ± 0.53	0.619 ± 0.025	0.869 ± 0.014
FedAvg	82.3 ± 0.47	0.623 ± 0.021	0.865 ± 0.012
FedProx	82.0 ± 0.41	0.615 ± 0.023	0.872 ± 0.011
FedNova	81.8 ± 0.52	0.621 ± 0.020	0.868 ± 0.015
FedGov-DP (ours)	81.5 ± 0.44	0.614 ± 0.022	0.867 ± 0.013
Local SGD	81.3 ± 0.58	0.607 ± 0.028	0.859 ± 0.017
DP-FedAvg	79.3 ± 0.51	0.581 ± 0.026	0.847 ± 0.018
Local Only	74.2 ± 1.31	0.518 ± 0.037	0.804 ± 0.026

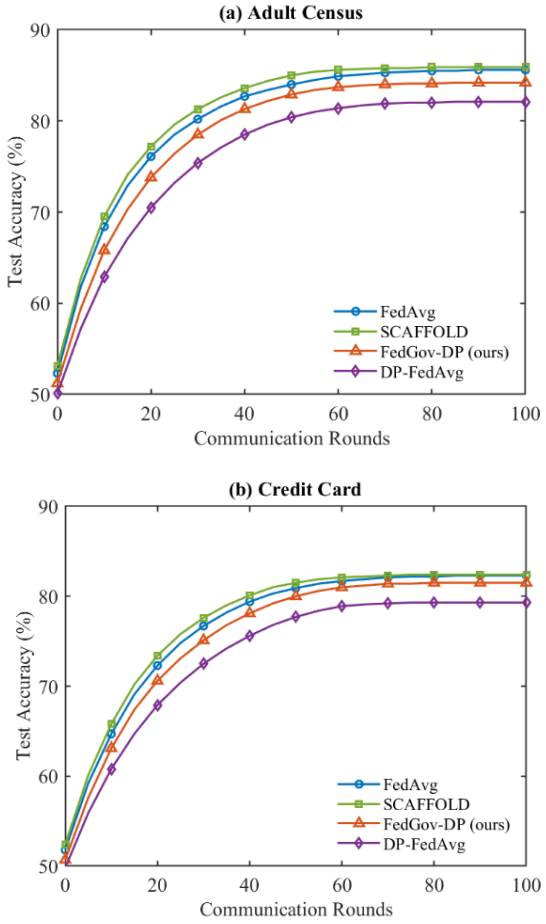

Figure 2. Convergence curve comparison of different methods on two datasets

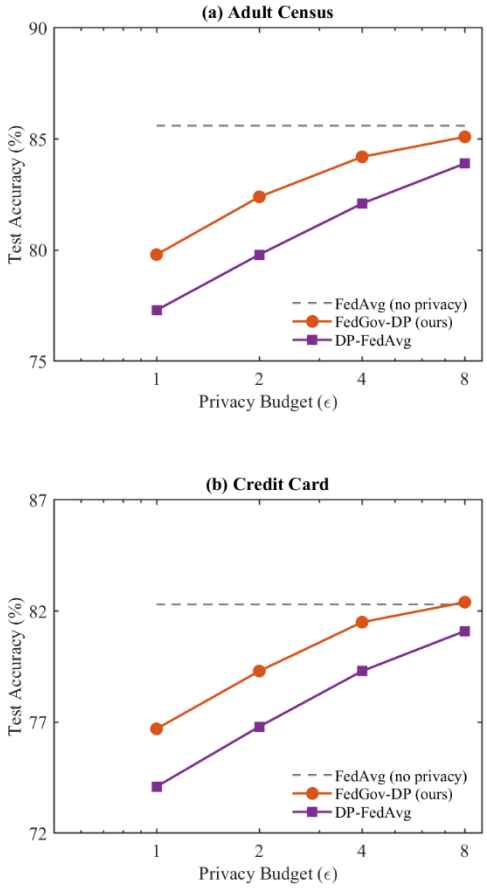

Figure 3. Trade-off curves between privacy budget ε and model accuracy

Figure 3(a) shows that in the Adult Census dataset, the accuracy of FedGov-DP steadily improves with the privacy budget to reach the FedAvg baseline while remaining better than DP-FedAvg in all cases. Figure 3(b) shows the same for the Credit Card dataset. The recommended value of the privacy budget ε∈[2,4] reflects the typical data protection requirements in public administration that follow the tenets of tiered government data protection, where ε≤2 ensures stronger data protection that is required for very sensitive data such as individual health and financial data that is subject to stringent requirements, while ε=4 ensures acceptable data protection that can be used in moderately sensitive aggregate data where analytical use is of primary importance.

The privacy-utility study presented above is a mathematical formulation of the tradeoff, but real-world privacy protection needs to resist empirical attacks. Membership Inference Attacks (MIA) is the approach to guess if the individual data point belongs to the training pattern or not, with success metrics tending to the random guess rate of 50% implying better privacy. The attack evaluation of membership inference follows the shadow model approach, where the attacking model has access to an auxiliary data set that is generated from the same distribution as the training data, but it can only access the prediction probability of the global model. The attacking model builds their own model to simulate the federated learning process in an effort to distinguish members from non-members according to their prediction confidence. The attack success rate is shown in Table 5.

Table 5. Membership inference attack success rates (%)

Method	Adult Census	Credit Card
FedAvg (no privacy)	67.3 ± 1.42	65.1 ± 1.58
DP-FedAvg (ε=4)	55.4 ± 0.93	54.1 ± 1.07
DP-FedAvg (ε=2)	52.7 ± 0.81	52.4 ± 0.94
FedGov-DP (ε=4)	53.8 ± 0.74	52.9 ± 0.86
FedGov-DP (ε=2)	51.2 ± 0.67	51.6 ± 0.79
Random Guessing	50.0	50.0

As shown in Table 5, the attack success rate without privacy protection in FedAvg is 67.3% and 65.1% for the two datasets, respectively, much higher than the random guess. This clearly verifies the information leakage. FedGov-DP ($\varepsilon = 4$) lowers the success rate to 53.8% and 52.9% respectively, performing better than DP-FedAvg (55.4% and 54.1%). For $\varepsilon = 2$, FedGov-DP lowers the success rate to 51.2% and 51.6%, which is close to the random guess rate and indicates that the dual privacy mechanism is more effective as a defense against membership inference attacks.

3.3 Validation of cross-departmental collaboration scenarios

The last two sections have verified the validity of federated collaboration efficacy and the accuracy of privacy protection. The following section centers on the characteristics in smart governance applications, including different numbers of collaborating departments and differences in their distributions because of functional differentiation. Simulations based on the Adult Census dataset (the 6-department case is more appropriate for collaboration analysis) aim at the flexibility of the designed methodology, FedGov-DP, for the scale of collaboration and the heterogeneity of the distribution. The data analysis of the effects of the number of participants involves testing four collaboration scales with K∈ 2,3,4,6. It seeks to determine the effects of marginal collaboration gains on collaboration scales. The experimental data is shown in Table 6.

Table 6. Impact of number of participating departments on model performance

Number of Departments (K)	Accuracy (%)	F1-Score	AUC-ROC	Marginal Gain
2	80.1 ± 0.61	0.647 ± 0.024	0.854 ± 0.016	—
3	82.3 ± 0.54	0.668 ± 0.019	0.871 ± 0.013	+2.2
4	83.5 ± 0.49	0.681 ± 0.022	0.879 ± 0.015	+1.2
6	84.2 ± 0.43	0.683 ± 0.017	0.891 ± 0.012	+0.7

From Table 6, the improvement in accuracy from 80.1% to 84.2% for K=2 to K=6 comes with a diminishing marginal increment of +2.2 percentage points for K=3, +1.2 percentage points for K=4, and +0.7 percentage points for K=6. The diminishing marginal increment can be explained by the law of diminishing marginal returns in which the marginal information derived from the additional departments will gradually diminish as the cumulative data set approaches the underlying population distribution, while the costs of collaboration will rise with the number of departments. Data heterogeneity is one of the key performance influencers of federated learning. When it comes to government applications, the Non-IID data distributions created due to functional departmentalization can lead to drifting and decreased performance of the global model. Heterogeneity is controlled in the experiments by the value of the Dirichlet

parameter α: for example, a value of $\alpha = 0.1$ reflects large heterogeneity (e.g., departments answering all inquiries for completely different populations), while a value of $\alpha = 5.0$ corresponds to near-IID (e.g., departments answering all inquiries for highly overlapping populations). Four group settings, with $\alpha \in \{0.1, 0.5, 1.0, 5.0\}$ are defined to span the entire range. Figure 4 presents the performance under different Non-IID intensities.

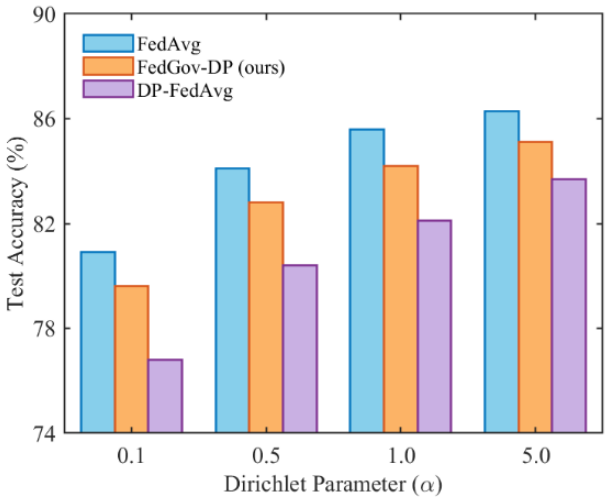

Figure 4. Impact of data heterogeneity (dirichlet parameter α) on model accuracy

Figure 4 results show that FedGov-DP reaches an accuracy of 85.1% for near-IID data ($\alpha = 5.0$), 84.2% for mild heterogeneity ($\alpha = 1.0$), 82.8% for moderate heterogeneity ($\alpha = 0.5$), and only 79.6% for under strong heterogeneity ($\alpha = 0.1$). The degradation of performance remains controllable (about 2 percentage points) in mild/moderate scenarios of heterogeneity. It shows strong adaptability to the common functional differentiation of departments in a government setting. Meanwhile, the approximate 5 percentage points degradation in strong scenarios indicate that the adaptability of the proposed method in extreme situations of Non-IID needs further improvement.

4. Discussion

To verify the effectiveness of the privacy-preserving federated learning framework for data collaboration across departments in smart governance, this study conducts systematic experiments, leading to three key findings. Federated learning obtains substantial data collaboration benefits even under the constraint of data sovereignty, with FedGov-DP enhancing performance by 10.7 and 7.3 percentage points compared with local-exclusive training, with a difference of no more than 2.6 percentage points compared with centralized training, covering the applicability of federated learning in the medical field [26], and this research further explores the feasibility of the "data usable but invisible" paradigm in public governance, indicating the successful implementation of federated learning. Effective defense against attacks from the differential privacy mechanism has been achieved, with FedGov-DP improving membership inference attack success rates from 67.3% to 53.8% for ε=4, and further to 51.2% for ε=2, in line with the simulation of privacy leakage [27]. Collaboration among multiple departments has reached the point of diminishing returns, with accuracy increased by 4.1 percentage points from 2 to 6 members, and the mechanism has good adaptability to moderately heterogeneous data [28]. From the perspective of theories, this proposed approach can bridge the theory and practice gap between the concept of machine learning in the context of distributed architectures and the realities of public administration by providing a systematic approach towards the collaboration of data among agencies to tap the maximum amount of intelligence, because this new paradigm of governance requires the technological architecture to be conducive to efficiencies as well as compliance.

Implications for practice may be drawn from three aspects. Technical recommendations are the application of federated learning in government data-sharing alternatives and technical standards and interfaces for the public sector, based on research findings regarding the application of technology in the digital transformation of the public sector [29]. Sensitivity levels must guide the setting of privacy parameters, with values set at ε ≤ 2, which are suitable for highly sensitive data, and ε = 4, which is suitable for general data use [30]. Pilots must choose the number of departments, between 3 and 5, depending on the complexity of inter-departmental organization coordination [31]. This work also recognizes some limitations. At the data level, the experiments are based on public tabular data sets that simulate government scenarios instead of real-world administrative data [32], and both data sets are binary classification problems with structured data. For multi-class problems, the FedGov- DP framework does not require any architectural changes since the privacy-preserving methods are based on the gradients that are independent of the number of classes in the outputs, unlike unstructured data types that are common in modern government activities such as images of documents and text. From the technical perspective, the existing framework is not able to handle Byzantine opponents that may provide cheating updates [33], while the secure aggregation algorithm is of quadratic complexity in communications [34]. The possible future work includes: testing the framework in real-life governmental pilots, adding the use of permissioned blockchain for creating immutable audit trails, and developing game theory-based incentive mechanisms to mitigate the free-rider problem in the voluntary collaboration between the departments.

5. Conclusion

Data silos hinder the growth of smart governance. The conventional approach to data sharing is difficult to achieve effectively because there are privacy regulations and trust issues between departments that hinder the process. Therefore, there is a compelling need for innovative technological solutions that make data useful while maintaining privacy. The research work presents FedGov, a privacy-preserving federated learning solution designed for smart governance to address the deficit in innovative technological solutions for data useful but not visible in smart governance. The solution for this is based on a three-tier architecture: data, computation, and coordination to accommodate collaboration between multiple departments with diverse data sources. The solution is based on FedGov-DP, which uses differential privacy along with secure aggregation to achieve the essence of data useful but invisible. By conducting systemic experimentations regarding the Adult Census and Credit Card public data sets based upon FedGov-DP, it has been proved that FedGov-DP achieves a prominent silo-breaking effect in government data silos; federated joint optimization increases accuracy by 7%~11% compared to

local optimization alone, whereas differences in accuracy are kept under 3% relative to fully centralized optimization; membership inference attacks are efficiently protected by Differential Privacy mechanisms in that attack success rates are decreased to a level close to random guessing; and privacy parameter ε=2 to ε=4 stands as a recommended setting in practice. The theoretical significance of this research is that it generalizes the scope of federated learning from commercial fields to public governance and aims to create a theoretical foundation to handle the special characteristics of data collaboration in public governance within hierarchical administrative structures and different levels of data sensitivities of different functional departments. Not only does this work open a novel technological way toward resolving data silos in government effectively, but it also provides a meaningful technical solution to facilitate the smart governance digital transformation based upon the designed framework of FedGov that has a great vision in stimulating further verification and optimization attempts in practical government contexts toward promoting the further application of privacy-preserving federated learning technology in practices related to public governance for high-quality digital government development in the future.

Ethical issue

The author is aware of and complies with best practices in publication ethics, specifically regarding authorship (avoidance of guest authorship), dual submission, manipulation of figures, competing interests, and compliance with research ethics policies. The author adheres to publication requirements that the submitted work is original and has not been published elsewhere.

Data availability statement

The manuscript contains all the data. However, more data will be available upon request from the authors.

Conflict of interest

The author declares no potential conflict of interest.

References

[1] J. R. Gil-Garcia, A. Guler, T. A. Pardo, and G. B. Burke, "Characterizing the importance of clarity of roles and responsibilities in government inter-organizational collaboration and information sharing initiatives," Government Information Quarterly, vol. 36, no. 4, p. 101393, 2019. https://doi.org/10.1016/j.giq.2019.101393

[2] C. Dong, J. Liu, and J. Mi, "How to enhance data sharing in digital government construction: A tripartite stochastic evolutionary game approach," Systems, vol. 11, no. 4, p. 212, 2023. https://doi.org/10.3390/systems11040212

[3] G. Hammerschmid, E. Palaric, M. Rackwitz, and K. Wegrich, "A shift in paradigm? Collaborative public administration in the context of national digitalization strategies," Governance, vol. 37, no. 2, pp. 411-430, 2024. https://doi.org/10.1111/gove.12778

[4] X. Zhang, "A more secure framework for open government data sharing based on federated learning," Government Information Quarterly, vol. 41, no. 4, p. 101981, 2024. https://doi.org/10.1016/j.giq.2024.101981

[5] L. Lyu, H. Yu, J. Zhao, and Q. Yang, "Threats to federated learning," in Federated Learning: Privacy and Incentive: Springer, 2020, pp. 3-16. https://doi.org/10.1007/978-3-030-63076-8_1

[6] Q. Li et al., "A survey on federated learning systems: Vision, hype and reality for data privacy and protection," IEEE Transactions on Knowledge and Data Engineering, vol. 35, no. 4, pp. 3347-3366, 2021. https://doi.org/10.1109/TKDE.2021.3124599

[7] C. Zhang, Y. Xie, H. Bai, B. Yu, W. Li, and Y. Gao, "A survey on federated learning," Knowledge-Based Systems, vol. 216, p. 106775, 2021. https://doi.org/10.1016/j.knosys.2021.106775

[8] V. Mothukuri, R. M. Parizi, S. Pouriyeh, Y. Huang, A. Dehghantanha, and G. Srivastava, "A survey on security and privacy of federated learning," Future Generation Computer Systems, vol. 115, pp. 619-640, 2021. https://doi.org/10.1016/j.future.2020.10.007

[9] X. Yin, Y. Zhu, and J. Hu, "A comprehensive survey of privacy-preserving federated learning: A taxonomy, review, and future directions," ACM Computing Surveys (CSUR), vol. 54, no. 6, pp. 1-36, 2021. https://doi.org/10.1145/3460427

[10] R. Gosselin, L. Vieu, F. Loukil, and A. Benoit, "Privacy and security in federated learning: A survey," Applied Sciences, vol. 12, no. 19, p. 9901, 2022. https://doi.org/10.3390/app12199901

[11] K. Hu, S. Gong, Q. Zhang, C. Seng, M. Xia, and S. Jiang, "An overview of implementing security and privacy in federated learning," Artificial intelligence review, vol. 57, no. 8, p. 204, 2024. https://doi.org/10.1007/s10462-024-10846-8

[12] S. Pandya et al., "Federated learning for smart cities: A comprehensive survey," Sustainable Energy Technologies and Assessments, vol. 55, p. 102987, 2023. https://doi.org/10.1016/j.seta.2022.102987

[13] Y. Y. Ghadi et al., "Integration of federated learning with IoT for smart cities applications, challenges, and solutions," PeerJ Computer Science, vol. 9, p. e1657, 2023. https://doi.org/10.7717/peerj-cs.1657

[14] S. Rani, A. Kataria, S. Kumar, and P. Tiwari, "Federated learning for secure IoMT-applications in smart healthcare systems: A comprehensive review," Knowledge-based systems, vol. 274, p. 110658, 2023. https://doi.org/10.1016/j.knosys.2023.110658

[15] S. P. Karimireddy, S. Kale, M. Mohri, S. Reddi, S. Stich, and A. T. Suresh, "Scaffold: Stochastic controlled averaging for federated learning," in International conference on machine learning, 2020: PMLR, pp. 5132-5143. Available: https://proceedings.mlr.press/v119/karimireddy20a.html

[16] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Artificial intelligence and statistics, 2017: PMLR, pp. 1273-1282. Available: https://proceedings.mlr.press/v54/mcmahan17a.html

[17] S. Zhang, W. Yuan, and H. Yin, "Comprehensive privacy analysis on federated recommender system

against attribute inference attacks," IEEE Transactions on Knowledge and Data Engineering, vol. 36, no. 3, pp. 987-999, 2023. https://doi.org/10.1109/TKDE.2023.3297040

[18] J. C. Jiang, B. Kantarci, S. Oktug, and T. Soyata, "Federated learning in smart city sensing: Challenges and opportunities," Sensors, vol. 20, no. 21, p. 6230, 2020. https://doi.org/10.3390/s20216230

[19] D. A. E. Acar, Y. Zhao, R. M. Navarro, M. Mattina, P. N. Whatmough, and V. Saligrama, "Federated learning based on dynamic regularization," arXiv preprint arXiv:2111.04263, 2021. Available: https://arxiv.org/abs/2111.04263

[20] K. Wei et al., "Federated learning with differential privacy: Algorithms and performance analysis," IEEE transactions on information forensics and security, vol. 15, pp. 3454-3469, 2020. https://doi.org/10.1109/TIFS.2020.2988575

[21] K. Bonawitz et al., "Practical secure aggregation for privacy-preserving machine learning," in proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, 2017, pp. 1175-1191. https://doi.org/10.1145/3133956.3133982

[22] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," Proceedings of Machine learning and systems, vol. 2, pp. 429-450, 2020. Available: https://proceedings.mlsys.org/paper_files/paper/2020/hash/38af86134b65d0f10fe33d30dd76442e-Abstract.html

[23] J. Wang, Q. Liu, H. Liang, G. Joshi, and H. V. Poor, "Tackling the objective inconsistency problem in heterogeneous federated optimization," Advances in neural information processing systems, vol. 33, pp. 7611-7623, 2020. Available: https://proceedings.neurips.cc/paper/2020/hash/564127c03caab942e503ee6f810f54fd-Abstract.html

[24] X. Wu, Y. Zhang, M. Shi, P. Li, R. Li, and N. N. Xiong, "An adaptive federated learning scheme with differential privacy preserving," Future Generation Computer Systems, vol. 127, pp. 362-372, 2022. https://doi.org/10.1016/j.future.2021.09.015

[25] C. Chen et al., "Trustworthy federated learning: privacy, security, and beyond," Knowledge and Information Systems, vol. 67, no. 3, pp. 2321-2356, 2025. https://doi.org/10.1007/s10115-024-02248-x

[26] N. Rieke et al., "The future of digital health with federated learning," NPJ digital medicine, vol. 3, no. 1, p. 119, 2020. https://doi.org/10.1038/s41746-020-00323-1

[27] M. Nasr, R. Shokri, and A. Houmansadr, "Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning," in 2019 IEEE symposium on security and privacy (SP), 2019: IEEE, pp. 739-753. https://doi.org/10.1109/SP.2019.00065

[28] Q. Li, Y. Diao, Q. Chen, and B. He, "Federated learning on non-iid data silos: An experimental study," in 2022 IEEE 38th international conference on data engineering (ICDE), 2022: IEEE, pp. 965-978. https://doi.org/10.1109/ICDE53745.2022.00077

[29] M. Gasco-Hernandez, J. R. Gil-Garcia, and L. F. Luna-Reyes, "Unpacking the role of technology, leadership, governance and collaborative capacities in inter-agency collaborations," Government Information Quarterly, vol. 39, no. 3, p. 101710, 2022. https://doi.org/10.1016/j.giq.2022.101710

[30] R. Geyer, T. Klein, and M. Nabi, "Differentially private federated learning: A client level perspective," arXiv preprint arXiv:1712.07557, 2017. Available: https://arxiv.org/abs/1712.07557

[31] X. Liu and L. Zheng, "Cross-departmental collaboration in one-stop service center for smart governance in China: Factors, strategies and effectiveness," Government Information Quarterly, vol. 35, no. 4, pp. S54-S60, 2018. https://doi.org/10.1016/j.giq.2015.12.001

[32] D. C. Nguyen, M. Ding, P. N. Pathirana, A. Seneviratne, J. Li, and H. V. Poor, "Federated learning for internet of things: A comprehensive survey," IEEE communications surveys & tutorials, vol. 23, no. 3, pp. 1622-1658, 2021. https://doi.org/10.1109/COMST.2021.3075439

[33] M. Fang, X. Cao, J. Jia, and N. Gong, "Local model poisoning attacks to {Byzantine-Robust} federated learning," in 29th USENIX security symposium (USENIX Security 20), 2020, pp. 1605-1622. Available: https://www.usenix.org/conference/usenixsecurity20/presentation/fang

[34] J. So, B. Güler, and A. S. Avestimehr, "Turbo-aggregate: Breaking the quadratic aggregation barrier in secure federated learning," IEEE Journal on Selected Areas in Information Theory, vol. 2, no. 1, pp. 479-489, 2021. https://doi.org/10.1109/JSAIT.2021.3054610