



International Journal of Information and Communication Technology

ISSN online: 1741-8070 - ISSN print: 1466-6642

<https://www.inderscience.com/ijict>

Optimising privacy protection for cross-border data flow based on federated learning

Feifei Niu

DOI: [10.1504/IJICT.2026.10076595](https://doi.org/10.1504/IJICT.2026.10076595)

Article History:

Received:	20 November 2025
Last revised:	16 December 2025
Accepted:	20 December 2025
Published online:	10 March 2026

Optimising privacy protection for cross-border data flow based on federated learning

Feifei Niu

School of Artificial Intelligence and Big Data,
Guangzhou Vocational University of Science and Technology,
Guangzhou, 510555, China
Email: nff3345@163.com

Abstract: As cross-border data flows continue to expand, protecting people's privacy has become a major barrier to international cooperation. This research provides a privacy-preserving optimisation technique based on federated learning to address it. In this framework, a cross-border dataset is produced first, and then five main modules are designed. Then, a neural network is used as the base model for federated training to train a global model without sending any data outside of the country. Last, experimental evaluations check how well the system works. The results show that the proposed method has an accuracy of 0.892 and a robustness of 0.912, meaning it works well even while processing many types of data and protecting privacy. This is far better than standard methods. This technology offers a safe and effective way to work together on data across borders.

Keywords: federated learning; cross-border data flow; privacy protection; multimodal data.

Reference to this paper should be made as follows: Niu, F. (2026) 'Optimising privacy protection for cross-border data flow based on federated learning', *Int. J. Information and Communication Technology*, Vol. 27, No. 17, pp.52–70.

Biographical notes: Feifei Niu is a Lecturer at the Guangzhou Vocational and Technical University of Science and Technology, China. She received her Bachelor's in Computer Science and Technology (2011), followed by a Master's in Engineering (Computer Application Technology) (2014) from the Henan Polytechnic University, China. She received her PhD in Information Technology from the SEGi University, Malaysia, in 2025. She has published one paper indexed by SCI and one paper indexed by EI. Her research interests encompass information, data privacy protection, and artificial intelligence.

1 Introduction

With the rapid advancement of global digitalisation, data has become a pivotal element in socio-economic development (Abbas and Zaman, 2024). Cross-border data flows, in particular, underpin international trade and collaborative efforts. Through data exchange, enterprises and organisations have significantly enhanced industrial synergy and innovation. However, alongside the explosive growth in data volumes and the increasing complexity of the global regulatory landscape, cross-border data flows have also exposed

a series of security and privacy risks. Federated learning (FL) offers a solution for privacy protection in cross-border data flows that balances efficiency with regulatory compliance.

Still, there are several problems that need to be solved before FL may be used in cross-border data circulation. On one hand, cross-border data shows typical non-independent and identically distributed (non-IID) traits, with big differences between participants in terms of data volume, sample kinds, and distribution patterns (Whitmore et al., 2025). This slows down the process of training the model and could possibly impair the performance of the global model. On the other, in big cross-border projects, the many people involved, and the complicated network conditions might cause problems with communication latency and bandwidth limits. Traditional FL protocols frequently need to sync parameters often, which not only cost more to communicate but also makes it more likely that time inconsistencies will happen.

In the global research community, companies like Google, Meta, and NVIDIA have made FL a key area for machine learning that protects privacy. Its uses are growing from mobile devices to many other fields, such as healthcare, transportation, and finance. In academics, Kairouz et al.'s systematic analysis of FL provides a thorough overview of communication optimisation, heterogeneity management, and security methods; Bonawitz et al.'s secure aggregation protocol has become an industry standard. Although these studies establish the theoretical foundation for FL, current research is inadequate for the particular context of cross-border data flow. Most of the present work on privacy computing is about making it operate better from a technological point of view. It does not take into account data compliance, international collaboration mechanisms, or transnational trust frameworks in a systematic way. As a result, it is necessary to create a FL framework that meets international privacy requirements from a cross-border data governance point of view, so that both technology innovation and regulatory compliance may be improved.

The main goals of this research are threefold: first, to systematically examine the technical challenges and privacy protection needs in cross-border data flows in order to define the relevant scenarios and limitations of FL in cross-border data governance; second, to suggest a FL optimisation strategy that combines privacy controls with convolutional neural networks (CNNs). Thirdly, validate the feasibility of the proposed approach through experimentation, assessing its performance in terms of privacy protection and model accuracy.

In transnational financial cooperation, banks and regulators might work together to train credit risk models employing FL to assess credit risk across borders. In international medical research, healthcare organisations from many countries could jointly create illness diagnosis models without disclosing patient data, thus reconciling privacy protection with scientific collaboration. This research also gives policymakers a technological framework they may use to help them implement international data governance guidelines.

In conclusion, FL presents a novel technical approach to safeguarding privacy in cross-border data flows; nonetheless, its optimisation and application in intricate real-world contexts necessitate thorough examination. The goal of this paper is to build a FL optimisation model and a framework for cross-border data flow so that we can systematically look at how well privacy protection mechanisms work, how scalable algorithmic performance is, and how easy it is to follow compliance rules. This will give new technical support for global data collaboration. The paper is organised like this:

Section 1 talks about the research background, why it is important, what the current situation is, and what the goals are. Section 2 goes into further detail about the theoretical and technical bases that are important, such as FL principles, privacy protection systems, and cross-border data regulatory regimes. Section 3 suggests a way to optimise cross-border data transfer using FL. Section 4 explains how to create experiments and analyse the data to make sure that the model works and protects privacy. Section 5 ends with a summary and a look ahead, pointing out key areas for future research.

In summary, this article talks about how global data governance is changing and how FL may help preserve privacy in cross-border data flows. This will lead to safe, efficient, and compliant ways for people to work together with data. The results not only help the creative use of privacy computing technologies, but they also help construct an open and secure system for the flow of data around the world.

2 Relevant technologies

2.1 Federated learning

FL is a new way to distribute machine learning that lets different groups work together to train models without having to store raw data in one place (Zhou, 2025). The main idea behind it is to move the focus of model training from centralising data to centralising models. Participating nodes store data locally and only provide model parameters or gradient updates to a central server or aggregation node. This method protects privacy while also optimising the global model. As a result, many people think of it as an important technology for AI that protects privacy in the future.

In a standard FL framework, there are three main parts:

- 1 Clients, which store data locally and run model training.
- 2 Server, which collects model parameters uploaded by clients and updates the global model.
- 3 Communication protocol, which sets the rules for sending, aggregating, and synchronising parameters.

FL training usually happens in several cycles, with each round including a local update phase and a global aggregation phase (Ji et al., 2024).

Let w be the global model parameters and K be the number of clients who are taking part. Each client k has a local dataset D_k that has n_k samples in it. The formal expression of FL's optimisation target is the minimisation of the subsequent global loss function:

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

$$n = \sum_{k=1}^K n_k \quad (2)$$

where n stands for the total number of client samples, and $F_k(w)$ stands for the local loss function for the k^{th} client. It is usually defined as:

$$F_k(w) = \frac{1}{n_k} \sum_{i=1}^{n_k} f_i(w, x_i, y_i) \quad (3)$$

where $f_i(\cdot)$ is the loss function for one sample, and (x_i, y_i) is a pair of training samples and labels. This formula shows that FL's goal for global optimisation is to do a weighted sum of the local optimisation results from all the nodes that are taking part, which gets close to the global optimum without having to centralise the data.

The federated averaging (FedAvg) algorithm is the most common one used in real life. The main steps in its workflow are as follows:

- 1 The server sends the current global model parameters w_t to all of the clients that were chosen.
- 2 Each client does many rounds of gradient descent updates on its own dataset, which gives them locally updated model parameters w_t^k .
- 3 Clients send their updated parameters to the server.
- 4 The server uses a weighted average based on sample size to combine these updates into the new global model w_{t+1} , which is calculated as follows:

$$w_{t+1} = \sum_{k=1}^K \frac{n_k}{n} w_t^k \quad (4)$$

This weighted averaging system makes sure that nodes with bigger sample sizes have more of an effect on the global model, while keeping the model stable across different types of data. The FedAvg method is widely used in mobile collaborative training, medical data sharing, and cross-institutional learning since it is easy to understand and does not require much discussion.

But FL's real-world use goes well beyond just the algorithm. The model updates themselves could still be reverse engineered to find sensitive information. Some common ways to improve things are:

- Differential privacy (DP): adding random noise to model update parameters so that individual data cannot be identified (Ponomareva et al., 2023).
- Secure multi-party computation (SMPC): allowing numerous parties to work together to get global outcomes without showing private inputs (Mamidala, 2021).
- Homomorphic encryption (HE): allows end-to-end secure training by supporting combined operations on encrypted model parameters.

There are many different forms of FL, depending on the link between the participants' data feature space and sample space, the system topology, and the application scenarios. Table 1 lists some common types, along with their traits, benefits, and drawbacks.

The table's classifications show the different privacy protection priorities and system design differences that FL needs to deal with in different application settings. Horizontal FL works well when the sample space is different, but the feature space is the same. For example, it can be used to predict a user's credit across numerous institutions. Vertical FL is used when the feature space is different, but the users are the same. For example, it can be used for cooperative data analysis between businesses and banks. Federated FL is

suitable when involved entities have markedly distinct data distributions while possessing shared knowledge transfer needs.

Table 1 FL characteristics under different data distributions and collaborative structures

<i>Type</i>	<i>Data distribution characteristics</i>	<i>Model training features</i>	<i>Key privacy focus</i>	<i>Limitations</i>
Horizontal FL	Same feature space, different sample space	Each party trains on the same feature dimensions; model aggregation through parameter averaging	Protection of user identifiers and behavioural data	Requires sample ID alignment; risk of duplicate samples
Vertical FL	Same sample space, different feature spaces	Secure alignment of different features across parties	Protection of feature attributes and labels	Complex sample alignment and cryptographic computation; high communication cost
Federated TL	Partial overlap in both samples and features	Leverages transfer learning to enable cross-domain knowledge sharing	Protection of model parameters and structural privacy	Complex training; requires layer-level mapping for transfer
Hierarchical FL	Multi-layer aggregation (edge-cloud hierarchy)	Local aggregation at edge nodes before uploading to the cloud	Privacy of communication links and edge nodes	Complex aggregation layer design; synchronisation challenges
Secure FL	Arbitrary data distribution	Integrates cryptographic and differential privacy mechanisms	Security of model parameters and aggregation process	High computational cost; reduced communication efficiency

From an optimisation point of view, the local training and global aggregation that happen in each FL iteration can be seen as an alternating optimisation process of local gradient descent and global averaging fusion (Chen et al., 2020). In particular, local updates can be formally stated as:

$$w_{t+1}^k = w_t - \eta \nabla F_k(w_t)$$

where η is the learning rate and $\nabla F_k(w_t)$ is the gradient of the k^{th} client based on the current model parameters.

Overall, FL has emerged as a crucial foundational technology in cross-border data circulation due to its combined advantages of distributed computing and privacy protection.

2.2 Cross-border data flow

Against the backdrop of the globalised digital economy, cross-border data flows have become a crucial indicator of national competitiveness and international cooperation

capabilities. However, the privacy breaches, sovereignty disputes, and compliance risks arising from the free movement of data have also become increasingly prominent.

The compliance framework for cross-border data circulation often includes encryption, access controls, data anonymisation, and intelligent algorithm security. Deep learning (DL) algorithm security and explainability have become key problems in cross-border data compliance frameworks. Explainable AI (XAI) and fairness-constrained learning allow models to communicate decision justifications, feature contributions, and risk explanations alongside their outputs (Zhang, 2025). For example, when credit scores are used across borders, an algorithm that cannot explain why an applicant from a certain country was turned down is at risk of breaking the rules. By using XAI technology, the model may give decision reasons that can be checked based on feature weights. This builds trust in the model by regulators and lays the groundwork for worldwide algorithm certification.

Table 2 Global cross-border data governance frameworks and technical security strategies

<i>Framework</i>	<i>Legal focus</i>	<i>Main challenges</i>	<i>Algorithmic/technical approaches</i>	<i>Application scenarios</i>
European Union (GDPR)	Personal data protection, privacy by design	Strict consent and auditing	Differential privacy, explainable AI, secure multi-party computation	Health analytics, finance, e-government
United States (CCPA, HIPAA, DPF)	Transparency, accountability, sectoral regulation	Fragmented compliance regime	Homomorphic encryption, blockchain-based audit, DL security	Smart healthcare, autonomous systems
China (PIPL, DSL)	Data localisation, sovereignty, security review	Localisation cost and legal complexity	Trusted execution environment, federated/transfer learning, blockchain traceability	E-commerce, credit scoring, IoT
Japan (APPI)	Consent, international transfer management	Third-party data control	Split learning, differential privacy	Cloud services, industrial data sharing
Singapore (PDPA)	Accountability, consent management	Limited enforcement	Privacy-preserving DL, access control	Cross-border enterprise data, fintech
South Korea (PIPA)	Explicit consent, data breach notification	Severe penalties for violations	Explainable AI, TEE-based computation	Digital ID systems, cross-border R&D

The system architecture for cross-border data flow security typically includes the following modules:

- 1 Encrypted computation and privacy protection modules that do secure computational tasks.
- 2 Trusted audit and traceability modules that make sure people are held accountable through blockchain or logging mechanisms.

- 3 Algorithm compliance adaptation layers that automatically change model training parameters and communication frequencies based on national policies.
- 4 An international trust certification system for cross-platform, cross-jurisdictional mutual recognition of security trust.

By working together in a natural way, these components provide a complete closed-loop system that covers everything from basic data security to high-level algorithmic compliance (Kioskli et al., 2025).

Generative models like generative adversarial networks (GANs) and variational autoencoders (VAEs) have been used a lot in research on compliance data modelling and privacy protection in the last few years. They have come up with new ways to move data across borders. The main idea behind these generative models is to understand the distribution of real data so that they may make virtual samples that are statistically similar to real data but leave out some sensitive information. This method keeps the data's analytical utility while lowering the possibility of privacy leaks. GANs and VAEs are better at preserving distributions than standard desensitisation or anonymisation methods. They accomplish privacy replacement while preserving the intricate structural links and high-dimensional feature correlations characteristic of the original data.

Table 2 summarises the main legislative frameworks, compliance issues, and associated intelligent algorithmic solutions for cross-border data governance in different nations or regions. This is done in a methodical way to compare the characteristics of cross-border data governance and the technical remedies that go with it.

3 Optimisation scheme for cross-border data flow based on FL

3.1 System design and architecture


This paper presents an optimal cross-border data circulation system utilising FL to tackle the difficulties of privacy protection and compliance in cross-border data flows. The goal of this technology is to make it possible for various organisations and countries to work together on data safely and train intelligent models. FL is the key algorithm in its architecture. At the same time, to make security and compliance even better, the system uses a DP mechanism during federated training. This guarantees that individual data contributions cannot be reverse engineered when parameters are uploaded, which meets data protection rules in many different places.

The system's general design includes five main parts: data preparation, federated training, privacy protection, model aggregation and optimisation, and system monitoring and compliance auditing, as shown in Figure 1.

The operational workflow of the whole system follows complete closed-loop logic. First, each node does its own data pre-processing and training for CNN. After that, the DP method makes the model parameters less clear, and then SMPC securely combines them. The weighted averaging process combines the global model into one, and the knowledge is sent back to the local nodes. At the same time, the blockchain module keeps track of operational logs and model versions during the whole process, making sure that the system can be audited and traced. This closed-loop approach protects data across borders while also making model training more efficient and improving performance

around the world. It is a workable technical solution for data sharing between multiple countries and institutions.

Figure 1 System architecture of cross-border data circulation solution (see online version for colours)



First, the data pre-processing module oversees cleaning, standardising, and getting features from the local data of each institution that is taking part. This makes sure that the data formats for model input are the same and gets rid of noise and outliers. This module uses a CNN feature extractor for unified feature encoding to make the model better at generalising. This maps raw data to a feature space with a fixed number of dimensions. This keeps important information safe while lowering the danger of exposing sensitive data.

The federated training module is the system's main computing unit. Each node in the system runs the training task for the CNN model on its own. This module uses CNN as a single model architecture to make sure that training and feature representation are the same. After calculating model gradients locally, nodes use the DP technique to inject noise before sending the changed parameters to the central aggregation server. DP is very important in this module since it protects people's privacy while yet making sure that model training is accurate and useful.

Thirdly, the privacy protection module encrypts and handles sensitive data safely while it is being sent and combined. This module uses SMPC to safely combine the gradients of local models from several nodes that are taking part. With SMPC, the server cannot get any specific data or settings from individual nodes, which protects data from end to end (Cai et al., 2023).

Fourthly, this module also incorporates model distillation techniques, feeding knowledge from the global model back to individual nodes to optimise local training

outcomes while reducing cross-border communication traffic, thereby achieving a balance between training efficiency and performance.

Finally, the module oversees and audits the entire framework. Only blockchain technology is used to record model changes, parameter uploads, training logs, and compliance inspections in this module. Blockchain's immutability creates trustworthy audit trails in cross-border multi-institutional systems, giving regulatory agencies and internal business audits with verifiable evidence of system operations' transparency and traceability.

In general, this system protects data privacy in cross-border settings by setting the technical selection for each module, which also improves the performance of the global model. This gives a strong base for improving the algorithm, testing its effectiveness across borders, and putting it to use in real life.

3.2 *Optimisation algorithm and implementation mechanism*

3.2.1 *Data pre-processing module*

The module is the first step in cross-border data circulation systems. It is also the basis for making sure that model training is of high quality and that cross-institutional collaboration goes smoothly. Because the nodes that are taking part come from different countries and institutions, their data is very different in terms of format, distribution, and scale. Applying directly to model training could lead to unstable training or a drop in overall model performance. The main job of the data pre-processing module is to clean, standardise, and encode the features of raw data from each node so that all the data is in the same format, outliers are handled, and cross-border usability is improved.

Data cleansing is the first step in pre-processing. It includes dealing with missing values, getting rid of duplicate records, finding outliers, and filtering out noise. In cross-border situations, differences in data collection tools, sample rates, and storage standards make it necessary to include domain knowledge in the cleaning process (Guamán et al., 2021). This makes sure that important information is kept in post-processing data while unnecessary information is removed. For example, interpolation or filtering can be used to fix missing scan segments or image noise in medical imaging data. On the other hand, financial transaction data may use anomaly detection algorithms to find and fix or remove records of transactions that are not normal. Only cleaned data can ensure that the model is stable and consistent when training on different nodes.

Second, standardising data is a necessary step for working together across institutions. The data scales across the nodes that are taking part may be very different. Training neural networks directly on such different data can cause the model to favour features with bigger magnitudes, which can lead to an imbalance in training. The technique uses Z-score standardisation to make all the data scales the same. This changes each feature value x_i into a standardised score with a mean of 0 and a standard deviation of 1 (Schober et al., 2021). Here is how it works:

$$x'_i = \frac{x_i - \mu}{\sigma} \quad (5)$$

where μ is the mean of the feature in the local dataset, σ is the standard deviation, and x'_i is the standardised feature values.

Additionally, the data pre-processing module is responsible for feature extraction and encoding. For example, convolutional procedures pull out spatial hierarchical features from picture data, one-dimensional convolutions find local patterns in structured tabular data and sliding window convolutions find dynamic features in time-series data (Zheng et al., 2024). After feature extraction, the data from all nodes keeps the same dimensionality and distribution, which makes it stable for gradient aggregation and global optimisation during federated training.

The module adds simple data augmentation or denoising methods to the feature space to improve data quality and model robustness even more. For example, random rotation or translation can be used on image data to make the training samples more varied. For tabular or time-series data, small changes can make lightweight synthetic samples that lessen the effect of single-node data biases on the global model.

3.2.2 Federal training module

This module uses CNN as its main model architecture to handle several sorts of data, such as photos, tables, and time-series features. Each participating node trains its CNN model locally using data that has been cleaned, standardised, and feature-encoded, while keeping the data stored locally without sending it across borders (Hossain et al., 2022). This way, data privacy is effectively protected.

Federated training has two main steps: training a model on a local level and combining models on a global level. During the local training phase, each node uses its own data to calculate gradients and change the parameters of the CNN model. Let w_k^t be the local model parameters of node k , which were modified in iteration t using gradient descent according to the following formula:

$$w_k^{t+1} = w_k^t - \eta \nabla L_k(w_k^t) \quad (6)$$

where η stands for the learning rate, and $\nabla L_k(w_k^t)$ stands for the gradient of the loss function that node k calculated on its own data.

The aggregation method employed in this study is weighted average aggregation, calculated using the following formula:

$$w^{t+1} = \sum_{k=1}^K \frac{n_k}{n} w_k^{t+1} \quad (7)$$

This formula weighs each node's model based on the size of its data, which makes sure that the global model balances data differences between nodes and improves the performance of the model in general. This weighted averaging not only handles data differences between nodes, but it also makes sure that the global model's parameter updates are stable, which means that the model is consistent and very accurate in environments with data from many sources and regions.

Overall, the federated training module forms a complete closed-loop through local CNN model training, gradient updates, and weighted aggregation, enabling cross-node data collaboration without exposing raw data (Sinha et al., 2025).

3.2.3 Privacy protection module

Within the SMPC framework, each node decomposes the local model parameter w_k into multiple secret shares (Tian et al., 2022). For instance, after splitting into three shares, these are denoted as $\langle w_k \rangle_1$, $\langle w_k \rangle_2$, and $\langle w_k \rangle_3$, and distributed to different computational parties. The fundamental partitioning of secret shares can be described using the following formula:

$$w_k = \langle w_k \rangle_1 + \langle w_k \rangle_2 + \langle w_k \rangle_3 \quad (8)$$

When performing global model aggregation, SMPC enables parties to complete weighted summation operations without restoring the original parameters. The aggregator obtains the final model parameters by executing distributed computations on the secret shares, with the computational process represented as follows:

$$w^{t+1} = \sum_{k=1}^K \frac{n_k}{N} (\langle w_k^{t+1} \rangle_1 + \langle w_k^{t+1} \rangle_2 + \langle w_k^{t+1} \rangle_3) \quad (9)$$

Each secret share is input into the aggregation protocol by distinct computational nodes, with the global server remaining entirely isolated from any actual parameters throughout the process, thereby achieving genuinely zero-knowledge aggregation (Oude Roelink et al., 2024). As SMPC operates on secret shares rather than raw data, the entire computation remains fully transparent to participants, significantly enhancing the system's trustworthiness in cross-border and cross-organisational collaborations.

3.2.4 Model aggregation and optimisation module

In situations where data flows across borders, this module not only optimises algorithms but also makes sure that the aggregation process is easy to understand, strong, and follows the rules. This makes it possible for high-risk industries (including finance, healthcare, and cross-border e-commerce) to train together effectively while following strict rules. This system uses FedAvg as its main aggregation technique. It also uses an adaptive weighting strategy and a local update assessment mechanism to make convergence faster and more stable in cross-border heterogeneous data situations (Haripriya et al., 2025).

At the foundational level, the aggregation module first depends on the conventional FedAvg algorithm to update the global model. Let's say there are K customers taking part in the t^{th} round of communication. Each client k does a few local iterations based on its local data volume n_k and uploads local model parameters w_t^k . The server then gets new global model parameters w_{t+1} by averaging the weights:

$$w_{t+1} = \sum_{k=1}^K \frac{n_k}{\sum_{j=1}^K n_j} w_t^k \quad (10)$$

This algorithm makes sure that clients with more data have more weight in global training. This improves the model's capacity to fit over the whole sample distribution. However, in cross-border data environments, sample distributions between countries or institutions are frequently significantly uneven, displaying marked skewness. For instance, one location largely showcasing positive examples while another primarily

comprising negative examples. This stops classical weighted aggregation from getting rid of all the biases that come from non-IID data (Kirchmeyer et al., 2021). This method builds on FedAvg by adding an update evaluation and dynamic weight adjustment mechanism to make the aggregation strategy more flexible when there are differences in the data distribution.

Before each aggregation round, the server uses the gradient direction and update magnitude that each client sends to figure out how credible the local update is for that client. Let $\Delta w_t^k = w_t^k - w_t$ be the amount that client k 's update changes the global model. We can then use the following normalised method to show the credibility weight:

$$\alpha_k = \frac{\|\Delta w_t^k\|}{\sum_{j=1}^K \|\Delta w_t^j\|} \quad (11)$$

When some clients have updated magnitudes that are much larger than others or orientations that are quite different from others, their weights are automatically lowered to keep local extreme distributions from messing with the model's steady convergence. In the final aggregation, both data weights and credibility weights are combined to create dual factor weighting (Blagojević et al., 2023). This makes the model stronger in contexts with data from several domains.

This module has a local momentum mechanism to cut down on communication rounds and make better use of bandwidth when coordinating across borders. In cross-border situations, differences in the quality of national networks, privacy concerns, and overhead system can all cause delays and discrepancies in the timing and size of model updates that clients upload.

3.2.5 System monitoring and compliance audit module

The module is the basis for all the rules and trust in the cross-border FL system. While the system is running, the server keeps an eye on the training status and behaviour patterns of each client. This includes things like the number of local iterations, the size of gradient updates, the time it takes for communication rounds to respond, the status of trusted hardware, and any unusual update characteristics, such as sudden updates, strange update directions, or frequent disconnections to avoid being audited. This module adds a compliance score to each training round so that we can see how compliant each client was. This makes it possible to automatically evaluate how each customer behaved throughout that round. Let the behavioural feature vector of client k during training round t be represented as x_t^k , encompassing attributes such as update norm, latency, TEE verification findings, and audit log integrity markers. The algorithm then uses a predefined audit weight vector a to figure out this client's compliance score.

$$C_t^k = a^\top x_t^k \quad (12)$$

This score is not used at the audit level to change the overall weightings; instead, it is used as a quantitative measure of compliance.

In short, the module is both technical protection and the foundation of governance that lets cross-border FL systems work steadily and sustainably across countries. This module makes sure that training procedures are safe, clear, and follow the rules by using real-time monitoring, automated scoring, log traceability, and policy rule enforcement. It

also provides a controlled, verifiable, and auditable technical framework for cross-border data sharing.

4 Experimental results and analyses

4.1 *Experimental data*

The system created by this research institute requires collaborative modelling between different countries and organisations. This means that building a dataset must meet many different needs. These include the capacity to work across domains, different levels of sensitivity, different data schemas, and detailed labelling frameworks. But for now, public databases do not always match the needs of compliance studies in cases when data is shared across borders. It is therefore indisputably demonstrated that no publicly available dataset currently meets the four criteria of cross-border properties, privacy categorisation, multimodal structure, and compliant annotation.

To tackle this issue, this study develops a synthetic yet genuinely organised cross-border multimodal compliance research dataset – the cross-border federated compliance dataset (CBFCD). Based on the healthcare, finance, and logistics industries, this dataset shows how different countries and institutions classify data sensitivity, share it, and give people access to it. The dataset includes structured tabular data, image data, and text data. It covers common cross-border application situations such multinational diagnostic collaboration, international banking risk assessment, and cross-border supply chain monitoring.

This study first sets data fields and sensitivity levels based on industry norms throughout construction. Statistical modelling and rule-based limitations were subsequently utilised to produce data samples with unique distributional characteristics, assuring the preservation of real-world statistical attributes while excluding any identifying personal information. At the same time, we added DP budget fields, cross-border restriction labels, and multi-institutional data partitioning strategies to the dataset to fulfil the needs of FL and privacy computing experiments (Kalejaiye et al., 2025). This makes it possible to utilise it in trials that look at how well distributed training works, how to reduce privacy issues, and how to improve compliance tactics across borders.

The final CBFCD dataset has three types of industry data, data from many countries and regions, sensitive attribute annotations, data sharing strategy labels, high-dimensional modal characteristics, and metadata descriptions that are all in the same format. This creates an experimental environment that is very easy to manage and repeat for future studies. Table 3 shows all the information in the dataset.

4.2 *Experimental procedure*

To investigate the impact of global iteration count on the robustness of FL systems, this experiment utilised the merged CBFCD dataset, treating all node data as a unified whole for training. Employing a CNN as the model architecture, combined with DP and SMPC, the global iteration count was adjusted while maintaining fixed local iteration rounds, learning rate, and batch size to observe changes in model performance. The number of iterations increased incrementally from 50 to 500 rounds, with each subsequent phase

adding 50 rounds (Browning et al., 2021). The model's convergence capability was primarily evaluated using accuracy metrics. Experimental results are presented in Figures 2 and 3.

Figure 2 Accuracy variation across different iteration counts (see online version for colours)





Figure 3 Loss rate variation across different iteration counts (see online version for colours)



During the initial training phase (50–150 iterations), owing to the limited number of global aggregations, the model was unable to fully learn the distributional characteristics of each client's data. Accuracy fluctuates markedly between 71% and 76%, while loss rapidly decreases from an initial value of 1.2 to approximately 0.5.

Table 3 CBFCD dataset overview

Category	Country/region	Num. samples	Data type	Key features (examples)	Cross-border restriction	Sharing strategy	Label type
Healthcare	China	5,200	Medical images (CT), structured clinical records	CT slices, age, diagnosis code, treatment history	Strict (GDPR-equivalent, localised storage required)	Encrypted model update only	Binary and multi-class
Healthcare	Germany	4,850	Medical images (MRI), lab test tables	MRI scans, biomarkers, physician notes statistics	GDPR strict	Differential privacy sharing	Multi-class
Healthcare	Singapore	3,600	Structured EMR	Vitals, medication records, admission type	Moderate	Aggregated statistics	Binary
Finance	USA	10,000	Structured financial transactions	Amount, merchant type, user behaviour patterns	PCL-DSS strict	Model parameter exchange	Fraud/not fraud
Finance	UK	8,200	Credit scoring records	Credit history, loan status, income bands	GDPR moderate	Anonymised feature sharing	Binary
Logistics	Japan	6,400	Time-series IoT sensor data	Temperature, GPS trajectory, vehicle load	Export-control mild	Federated model sync	Regression
Logistics	Netherlands	5,700	Textual shipment records	Route descriptions, customs notes	Low	Raw-text embedding exchange	Multi-class

The key phase between 200 and 350 training cycles optimises system performance. With greatly reduced fluctuation amplitude, accuracy stabilises at 84%–86%. Loss steadily converges at 0.2 and 0.3. Since parameter updates are stabilising, the loss curve's consistent drop implies the model has captured crucial data trends.


When the number of iterations was increased to 400–500 rounds, the system reached a point of diminishing returns. The accuracy only went up little, to about 0.87, which is less than a 1% gain. The loss function, on the other hand, was stable in a range of 0.15–0.2. At this point, the performance improvements from more training could no longer make up for the steady rise in communication and computing costs.

The CBFGD dataset has a lot of different features; therefore the model needs to go through a lot of initial iterations to understand the statistical patterns of each one. Tests show that slightly increasing the number of iterations not only makes the model better at generalising on fused data, but it also makes the system more fault tolerant when there are gradient perturbations and strange nodes. We suggest setting 250–350 global iterations for FL tasks of similar size, as this strikes a good balance between convergence efficacy and resource use. This makes sure that the model converges enough while still being trained in an acceptable amount of time.

In the previous experiment, we looked at how global iteration counts affected the accuracy and resilience of models across the fused CBFGD dataset. We found that the best range was 250 to 350 iterations. This experiment further investigates performance differences between other DL models and the suggested cross-border privacy-preserving optimised FL system on integrated data, thoroughly assessing the system's benefits. We chose six model architectures to test: CNN, recurrent neural network (RNN), long short-term memory (LSTM), gated recurrent unit (GRU), transformer, and fully connected neural network (FCN).

We evaluate the system's performance using accuracy, F1-score, and robustness score (Chicco and Jurman, 2020). Experimental results are shown in Figure 4.

Figure 4 Results of comparative experiments between different models (see online version for colours)



All three measures show that the proposed approach outperforms standard models: accuracy of 0.892, F1-score of 0.879, and robustness score of 0.912. It has many advantages in multimodal fusion and cross-border privacy protection. In contrast, the standard models have accuracy scores from 0.776 (FCN) to 0.861 (CNN) and similar gradient distributions for F1-score and robustness metrics. Model adaptability varies by feature type.

Further investigation shows CNNs excel at handling fused data with accuracy 0.861, F1-score 0.847, and robustness 0.884. Transformers have higher robustness (0.872) than CNNs but lower accuracy (0.842). RNN and LSTM perform well in sequence feature processing but worse in multimodal data fusion than CNN and transformer. Recurrent structures have low adaptation to complicated picture or structured features with robustness indices of 0.828 and 0.835.

FCN scores the lowest in accuracy (0.776) and robustness (0.772), whereas GRU scores 0.802 and 0.798. The proposed system degrades performance by 0.02–0.03 under 5% random gradient perturbations, the lowest among models. This reduces less than CNN and transformer.

Overall, the FL system surpasses solo DL models in performance and robustness.

5 Conclusions

This study endeavours to address privacy protection challenges in cross-border data circulation through FL technology, innovatively integrating techniques such as DP, SMPC, and CNN. Experimental results demonstrate the system's outstanding performance in multimodal data scenarios, achieving not only effective privacy protection but also surpassing traditional models in key metrics including accuracy, F1-score, and robustness. This provides a reliable technical solution for secure cross-border data circulation.

However, this study has drawbacks. The datasets and experimental setups are small and simplistic, failing to account for complicated network situations. Future study will expand datasets, introduce adaptive aggregation algorithms, and validate the system's practicality in real-world operating scenarios to overcome these challenges. This will help build a more complete cross-border data circulation system.

Declarations

All authors declare that they have no conflicts of interest.

References

- Abbas, S.A. and Zaman, A. (2024) 'Does digitalisation help achieve (selected) socio-economic SDGs? Evidence from emerging economies', *Sustainable Development*, Vol. 32, No. 6, pp.6088–6103.
- Blagojević, B., Nordstroem, E-M. and Lindroos, O. (2023) 'A framework for defining weights of decision makers in group decision-making, using consistency between different multicriteria weighting methods', *International Journal of Forest Engineering*, Vol. 34, No. 2, pp.130–142.

- Browning, B.L., Tian, X., Zhou, Y. and Browning, S.R. (2021) 'Fast two-stage phasing of large-scale sequence data', *The American Journal of Human Genetics*, Vol. 108, No. 10, pp.1880–1890.
- Cai, J., Zhou, Z., Huang, Z., Dai, W. and Yu, F.R. (2023) 'Privacy-preserving deployment mechanism for service function chains across multiple domains', *IEEE Transactions on Network and Service Management*, Vol. 21, No. 1, pp.1241–1256.
- Chen, M., Poor, H.V., Saad, W. and Cui, S. (2020) 'Convergence time optimization for federated learning over wireless networks', *IEEE Transactions on Wireless Communications*, Vol. 20, No. 4, pp.2457–2471.
- Chicco, D. and Jurman, G. (2020) 'The advantages of the Matthews correlation coefficient (MCC) over F1 score and accuracy in binary classification evaluation', *BMC Genomics*, Vol. 21, No. 1, p.6.
- Guamán, D.S., Del Alamo, J.M. and Caiza, J.C. (2021) 'GDPR compliance assessment for cross-border personal data transfers in android apps', *IEEE Access*, Vol. 9, pp.15961–15982.
- Haripriya, R., Khare, N., Pandey, M. and Biswas, S. (2025) 'A privacy-enhanced framework for collaborative big data analysis in healthcare using adaptive federated learning aggregation', *Journal of Big Data*, Vol. 12, No. 1, p.113.
- Hossain, P.S., Chakrabarty, A., Kim, K. and Piran, M.J. (2022) 'Multi-label extreme learning machine (MLELMs) for bangla regional speech recognition', *Applied Sciences*, Vol. 12, No. 11, p.5463.
- Ji, S., Tan, Y., Saravirta, T., Yang, Z., Liu, Y., Vasankari, L., Pan, S., Long, G. and Walid, A. (2024) 'Emerging trends in federated learning: from model fusion to federated x learning', *International Journal of Machine Learning and Cybernetics*, Vol. 15, No. 9, pp.3769–3790.
- Kalejaiye, A.N., Shallom, K. and Chukwuani, E.N. (2025) 'Implementing federated learning with privacy-preserving encryption to secure patient-derived imaging and sequencing data from cyber intrusions', *International Journal of Scientific Research Archives*, Vol. 16, No. 1, pp.1126–1145.
- Kioskli, K., Grigoriou, E., Islam, S., Yiorkas, A.M., Christofi, L. and Mouratidis, H. (2025) 'A risk and conformity assessment framework to ensure security and resilience of healthcare systems and medical supply chain', *International Journal of Information Security*, Vol. 24, No. 2, pp.1–28.
- Kirchmeyer, M., Gallinari, P., Rakotomamonjy, A. and Mantrach, A. (2021) 'Unsupervised domain adaptation with non-stochastic missing data', *Data Mining and Knowledge Discovery*, Vol. 35, No. 6, pp.2714–2755.
- Mamidala, V. (2021) 'Enhanced security in cloud computing using secure multi-party computation (SMPC)', *International Journal of Computer Science and Engineering (IJCSE)*, Vol. 10, No. 2, pp.59–72.
- Oude Roelink, B., El-Hajj, M. and Sarmah, D. (2024) 'Systematic review: comparing zk-SNARK, zk-STARK, and bulletproof protocols for privacy-preserving authentication', *Security and Privacy*, Vol. 7, No. 5, p.e401.
- Ponomareva, N., Hazimeh, H., Kurakin, A., Xu, Z., Denison, C., McMahan, H.B., Vassilvitskii, S., Chien, S. and Thakurta, A.G. (2023) 'How to dp-fy ml: a practical guide to machine learning with differential privacy', *Journal of Artificial Intelligence Research*, Vol. 77, pp.1113–1201.
- Schober, P., Mascha, E.J. and Vetter, T.R. (2021) 'Statistics from A (agreement) to Z (z score): a guide to interpreting common measures of association, agreement, diagnostic accuracy, effect size, heterogeneity, and reliability in medical research', *Anesthesia & Analgesia*, Vol. 133, No. 6, pp.1633–1641.
- Sinha, P., Sahu, D., Prakash, S., Rathore, R.S., Dixit, P., Pandey, V.K. and Hunko, I. (2025) 'An efficient data driven framework for intrusion detection in wireless sensor networks using deep learning', *Scientific Reports*, Vol. 15, No. 1, pp.1–25.
- Tian, N., Guo, Q., Sun, H. and Zhou, X. (2022) 'Fully privacy-preserving distributed optimization in power systems based on secret sharing', *iEnergy*, Vol. 1, No. 3, pp.351–362.

- Whitmore, J., Mehra, P., Yang, J. and Linford, E. (2025) 'Privacy preserving risk modeling across financial institutions via federated learning with adaptive optimization', *Frontiers in Artificial Intelligence Research*, Vol. 2, No. 1, pp.35–43.
- Zhang, G. (2025) 'Transformer-based AI framework for optimising English teaching evaluation strategies: a data-driven and explainable approach', *International Journal of Information and Communication Technology*, Vol. 26, No. 9, pp.107–127.
- Zheng, K., Wang, J., Chen, Y., Jiang, R. and Wang, W. (2024) 'DDTCN: decomposed dimension time-domain convolutional neural network along spatial dimensions for multiple long-term series forecasting', *Applied Intelligence*, Vol. 54, No. 8, pp.6606–6623.
- Zhou, B. (2025) 'A data compliance sharing algorithm for intelligent connected vehicles empowered by federated learning', *International Journal of Information and Communication Technology*, Vol. 26, No. 36, pp.42–57.