Privacy-Preserving AI Development Using Federated Learning: A Case Study of RSU

Khanat Kruthkul*, Chetneti Srisa-An, and Sumana Kasemsawasdi

College of Digital Innovation Technology, Rangsit University, Pathum Thani 12000, Thailand
*Corresponding author, E-mail: khanat.k67@rsu.ac.th

Abstract

Artificial Intelligence (AI) development currently faces significant data privacy bottlenecks due to stringent regulations such as Thailand's PDPA and the GDPR. Traditional centralized learning systems, which gather raw data into a single repository, increase the risk of data leaks and privacy violations. This paper investigates the feasibility of transitioning Rangsit University's (RSU) data environment from centralized servers to decentralized systems using Federated Learning (FL). By adopting the "model to data" principle, FL enables collaborative model training across distributed departments while keeping sensitive personal data localized and secure. The methodology integrates the FedAvg algorithm with Differential Privacy (DP) mechanisms to mitigate risks such as gradient leakage and membership inference attacks. Experimental results demonstrate that while centralized learning achieved an accuracy of 98.50%, the federated system maintained a high accuracy of 96.20%. This study concludes that the marginal loss in accuracy represents a worthwhile trade-off for the substantial gains in data privacy and reduced network bandwidth. These findings establish a framework for AI development that ensures both high performance and robust regulatory compliance.

***Keywords*:** *federated learning, differential privacy, PDPA, FedAvg, Non-IID data, data privacy, re-identification attacks*

1. Introduction

RSU has been implemented in a centralized data environment. The problem is how RSU can share data with various departments without violating data privacy. In addition, there is data used in teaching and learning that may contain personal information such as real names and surnames. To study the feasibility of reducing the risk of sharing personal data between organizations, the FL concept was developed.

The primary goal of Privacy-Preserving AI Development Using Federated Learning (FL) is to enable collaborative AI model training across distributed data sources while ensuring strong privacy protection and regulatory compliance. Teaching, learning, and general office work at Rangsit University all require data sharing, but this process is currently more complicated than before due to new laws, such as GDPR/PDPA, which prohibit the sharing of personal data without permission. Federated Learning (FL) is based on the concept of not sending data for central processing, but rather sending the model to be processed where the data is located. This prevents data leakage and eliminates data privacy risks because personal data is never sent outside the organization. This research attempts to explore the feasibility of implementing this concept in universities, as all departments possess personal data. Sending personal data to a centralized server for processing is more difficult than processing it at the organization and then simply sending the results to a central authority.

Federated Learning (FL) (Li et al., 2020; Zhang et al., 2021) is therefore proposed as a new approach that shifts from "data to model" to "model to data," reducing security risks and improving computing efficiency at edge devices. Federated Learning (FL) is recognized as a technology that enables AI development to effectively comply with personal data protection laws such as GDPR (and Thailand's PDPA). Its key principle is "Bring the code to the data, not the data to the code," or in other words, Federated Learning is a "solution" that allows us to achieve AI that is just as intelligent (or more intelligent) while keeping personal data secure.

Machine learning/AI relies on large amounts of data for training and model building. A major problem today is that the data used is often personal and sensitive, such as health history or location

information. This data is vulnerable to attacks such as inference attacks, where malicious actors attempt to find specific information within the training dataset. This is compounded by increasingly stringent data protection laws (such as GDPR or PDPA).

Federated Learning (FL) is recognized as a technology that effectively enables AI development to comply with personal data protection laws such as GDPR (and Thailand's PDPA), based on the key principle of "Bring the code to the data, not the data to the code." For example:

• Healthcare: Hospital A has cancer patient data, and Hospital B also has it, but laws (such as PDPA or HIPAA) prohibit data sharing. FL allows both hospitals to collaborate on creating the most accurate AI for cancer detection without leaking patient data.

• Cybersecurity: Multiple banks can jointly train fraud detection systems without disclosing customer transaction data to competitors.

• Reduced Bandwidth Burden: Instead of sending terabytes of data files across networks, only small model parameter files are sent.

Furthermore, Article 25 of the GDPR requires systems to be designed to protect data from the outset. Federated Learning is designed to be decentralized from the ground up, reducing the risk of a single point of failure. If a central server is attacked, hackers would only obtain the untrained model, but not individual user data such as names, addresses, or medical history.

1.1 Privacy Laws in Thailand

Thailand enacted the PDPA (Personal Data Protection Act), which came into full effect on June 1, 2022, making the collection, use, and sharing of personal data more complex than in the past.

Rangsit University prioritizes the protection of personal data for its personnel and students by establishing a dedicated Personal Data Protection Information System. Furthermore, the University's ROPA (Record of Processing Activities) process has been implemented for relevant staff. The ROPA (Huth et al., 2019; Mironeanu & Aflori, 2021) analysis found that personal data is scattered across various departments, such as faculties, which store their own data and documents. Therefore, data management needs to change from the previous system, which relied on a centralized server for important university data.

A feasibility study on sharing personal data revealed that it was not feasible because all faculties and departments require personal data for their operations. Therefore, the Federated Learning (FL) concept was chosen, which allows for data storage while processing the data by sending code to the relevant departments.

1.2 Privacy-Preserving Mechanisms

Differential Privacy (DP) is used in Federated Learning (FL) to enhance the protection of personal data. Although FL does not transmit raw data from the device, model updates or gradients can still leak information. Therefore, DP is used to mitigate this risk.

Differential privacy quantifies the guarantee that outputs from a learning system are insensitive to changes in any individual record. In FL, noise can be added to model updates before sharing, protecting sensitive information.

1.3 Hacker Attacks in Federated Learning Environment

Although Federated Learning (FL) has the advantage of not sending "raw data" out of the device, FL alone cannot guarantee 100% privacy. In FL, devices send "gradients," or model enhancement values, to a central server. Researchers found that hackers could use reconstruction attacks to trace these gradient values back to the user's raw "image" or "text."

Differential Privacy (DP) involves adding noise to communicated values to ensure that no one can distinguish how "an individual's data" affects the model. DP can help protect against hackers as follows:

1.3.1 Prevent gradient leakage:

DP adds statistical noise to the gradients before outputting them, making it difficult for hackers to clearly recover the original data. However, the model can still learn the overall statistical picture.

Proceedings of RSU International Research Conference (RSUCON-2026)

1.3.2 Preventing Membership Inference Attacks:

Membership Inference Attacks (Hu et al., 2022) are situations where hackers can guess who is in a group. DP provides a mathematical guarantee that the results of a model with Mr. A's data and a model without Mr. A's data will be "almost indistinguishable," making it impossible to identify who participated.

1.3.3 Prevent memorization

Large models (such as LLMs) often "remember" recurring or distinctive data, such as identification numbers or addresses (Wei et al., 2025). If a model remembers these values, subsequent users might randomly encounter this data. The DP process limits the influence of a single data point to prevent it from overly affecting the model, allowing the model to learn only "patterns" rather than "memorizing" individual data points.

1.4 Research Objectives

The purpose of this research is to study the feasibility of applying FL to reduce the risks of sharing information between departments as follows:

1.4.1 Develop a privacy-preserving Federated Learning (FL) framework.

1) Design a FL architecture that does not require centralizing raw data.
2) Integrate Differential Privacy (DP) and Secure Aggregation (SA) mechanisms.
3) Reduce the risk of data leaks through model updates.

1.4.2 Establish guidelines for AI development that comply with data protection laws.

1) Support compliance with GDPR, HIPAA, and PDPA.
2) Promote the concept of Ethical AI.
3) Increase confidence in the use of decentralized AI.

The research paper's structure consists of six sections: Chapter 1 introduces the research topic and presents the problem statement, Chapter 2 reviews and synthesizes relevant articles, Chapter 3 details the methodology employed, Chapter 4 presents and interprets the experimental results, Chapter 5 analyzes and discusses the experimental findings, and Chapter 6 summarizes the main conclusions and implications of the research.

2. Literature Review

Xu et al. (2021) reviewed the Privacy-Preserving ML (PPML) approach, focusing on the problems of handling large amounts of data that are susceptible to leaks, such as membership inference attacks or model inversion attacks. The authors propose a Phase-Guarantee-Utility (PGU) triad to evaluate PPML solutions based on phase (training/inference), privacy guarantees (e.g., differential privacy), and model utility. This article shows that designing a PPML system is not just about choosing one technique, but about considering the entire ML process and finding the right balance between privacy, accuracy, and resources.

Zhou et al. (2021) proposed applying Federated Learning (FL) to enhance the Industrial Internet of Things (IIoT) because IIoT has a huge volume and is often highly sensitive (Sensitive Data). It is very difficult to integrate data into a central cloud. Federated Learning (FL) is therefore a better solution because it does not need to send data to a central processing unit.

Ding et al. (2022) proposed a visionary perspective on the application of Federated Learning (FL) in large-scale federated systems. The paper emphasizes that FL is not solely for prediction but must also encompass detection and handling of highly heterogeneous data, with the goal of maintaining data privacy and reducing the cost of developing models on edge devices.

Najafi et al. (2024) demonstrated a solution to the main challenge in Federated Learning (FL), namely the Non-IID (Non-Independent and Identically Distributed) problem, where data on each device is very different, which causes the global model to be less efficient and more difficult to train.

Chatsuwan et al. (2023) proposed a privacy policy scoring model to systematically and quantitatively assess compliance with Thailand's Personal Data Protection Act (PDPA). They studied the compliance of

384 Thai SMEs using stratified random sampling during a two-year deferral period for the PDPA's enforcement. They provided recommendations for improving compliance, and the model is adaptable to the data protection laws of other countries.

Sweeney (2012) focused on the concept of k-anonymity to protect personal data in datasets intended for publication. She demonstrated that simply deleting identifying information (such as names or addresses) is insufficient to prevent identity theft, as malicious actors can link data using quasi-identifiers. This work presented the use of generalization and suppression to make data k-anonymous, described potential attack methods, and proposed best practices to mitigate vulnerabilities. The research showed that only a zip code, gender, and date of birth can identify personal information in the United States. Quasi-identifiers are defined as the minimum groups or sets of attributes that can identify an individual. Sweeney revealed that data privacy issues arise from combining voter registration data and public data sets to reconstruct original personal information.

Korolova et al. (2009) presented privacy protection techniques for publishing search queries and click data. Protecting the privacy of user search queries and click behavior is a crucial consideration, as this data can be sensitive and reveal personal preferences or identifiable patterns.

Abadi et al. (2016) focuses on training Deep Neural Network models using the principle of Differential Privacy (DP) to protect personal data in training datasets. The authors proposes the Differentially Private Stochastic Gradient Descent (DP-SGD) algorithm, which controls the sensitivity of each training data point through gradient clipping and noise addition using a Gaussian mechanism. He developed a tool called Moments Accountant to calculate privacy loss more accurately and efficiently than the traditional composition method.

Machanavajjhala et al. (2011) focuses on the trade-off between accuracy and privacy in personalized social recommendations. Personalized recommendations, such as those provided by social media platforms, rely on analyzing users' preferences and behavior to recommend relevant items or connections. However, in doing so, privacy concerns arise when the analysis of personal data may reveal sensitive information about individuals.

Monteiro et al. (2024) focuses on developing design patterns for data privacy in the context of GDPR compliance, emphasizing user privacy. It proposes four new models: Generalization, Hierarchical Generalization, Suppress Outliers, and Relocate Outliers, to strike a balance between data privacy and data usefulness.

Narayanan and Shmatikov (2008) proposed a novel statistical attack method that can identify users from large anonymized datasets, even with security measures in place. Their research focused on the Netflix Prize dataset, which contains movie ratings from over 500,000 users, demonstrating that attackers with only minimal information such as the few movies users have watched, along with ratings and approximate dates could accurately identify users in the database. The algorithm's strength lies in its ability to handle perturbations and background data errors.

Terrovitis et al. (2008) focuses on the dissemination of personal movement data collected from RFID or credit cards. Malicious actors can use portions of this movement data to identify other locations the cardholder has visited. Therefore, he proposes a technique to suppress (delete) certain data points from the movement path, preventing the identification of any remaining locations beyond a predetermined probability. He also designs a greedy algorithm to efficiently select points for deletion, considering both the security and quality of the remaining data.

Hu et al. (2022) presented Membership Inference Attacks (MIAs) on Machine Learning (ML) models, published in the 2022 ACM Computing Surveys. MIAs aim to investigate whether a sample dataset is used to train a target model, which could lead to privacy violations, such as inferring diseases from medical information.

Gong et al. (2022) state that in an era where personal data processing is being moved from our machines to servers of service providers (such as Google Cloud, AWS), this research paper explores personal data inference attacks on cloud models, focusing on models, technologies, and future research directions such as membership inference, model inversion, and property inference in the cloud context.

Ye et al. (2022) proposes a method to protect against data inference attacks using differential privacy (DP) by adjusting the confidence score vector from a machine learning model to handle both membership inference attacks (MIA) and model inversion attacks (MIA) by adjusting only one parameter: the privacy budget ε. This method normalizes the vector to obscure membership data or reconstructed data without changing the score order, thus preserving classification accuracy. This method is suitable for models that have already been trained and do not need to be retrained.

Garfinkel et al. (2019) presented a problem of database reconstruction attacks from public statistics tables, using a small census block example to demonstrate that microdata can be accurately reconstructed.

Mammen (2021) conducted a survey on Federated Learning (FL), as recommended by Google in 2016, summarizing opportunities in healthcare, finance, transportation, and NLP, along with training and security challenges.

3. Methodology

This study proposes a privacy-preserving federated learning framework that integrates:

1. Federated Averaging (FedAvg)
2. Differential Privacy (DP)
3. Secure Aggregation (SA)

The objective is to evaluate the trade-off between model performance and privacy guarantees under decentralized training conditions.

Workflow:

1. The server initializes global model parameters w_0.
2. At round t, a subset of clients $S_t \subseteq K$ is selected.
3. Each client performs local training using its private dataset.
4. Local updates are privatized using differential privacy.
5. Encrypted updates are transmitted using secure aggregation.
6. The server aggregates updates to obtain the new global model.

4. Experiment

This research uses PyTorch for several reasons. For example, much research in Fluorescent Dynamics (FL) involves solving data heterogeneity problems, which require tweaking loss functions or performing unusual gradient descent processes. PyTorch makes it easier to perform custom loss functions or intervene with variables than other frameworks that focus on static/graph-based operations. To simulate heterogeneous data distribution (non-IID), experiments are conducted using: MNIST (image classification) and CIFAR-10. Data are partitioned among clients using an IID distribution and a non-IID label skew distribution. The experimental environment was implemented using Python 3.9 and PyTorch 2.x. PyTorch was selected due to its dynamic computation graph, ease of model parameter manipulation, and strong support for GPU acceleration. These features are particularly beneficial for Federated Learning implementation, where local model updates and parameter aggregation are essential.

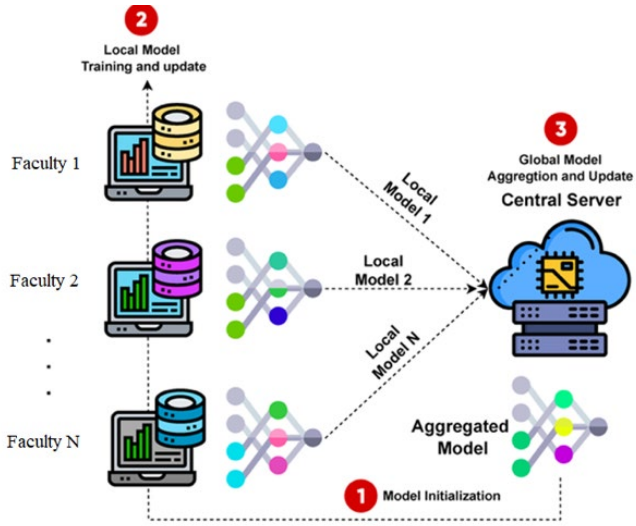

Figure 1 System Architecture

In the experiment, we compared the performance between Centralized Learning and Federated Learning using the MNIST and CIFAR-10 datasets:

- Setup: 100 clients were used, simulating a non-IID dataset (non-uniformly distributed data).
- Metrics: Accuracy and Communication Overhead were measured.
- Results: The experimental results showed that Federated Learning (FL) could achieve accuracy close to that of a Centralized Model (difference of less than 3-5%), but its key advantage is that 100% of the raw data was never sent from the user's machine.

Table 1 Performance Comparison between Centralized and Federated Learning on MNIST Dataset

Method	Accuracy	Data Privacy	Network Bandwidth
Centralized	98.50%	Low	High (Upload Data)
Federated (FedAvg)	96.20%	High	Low (Upload Weights)

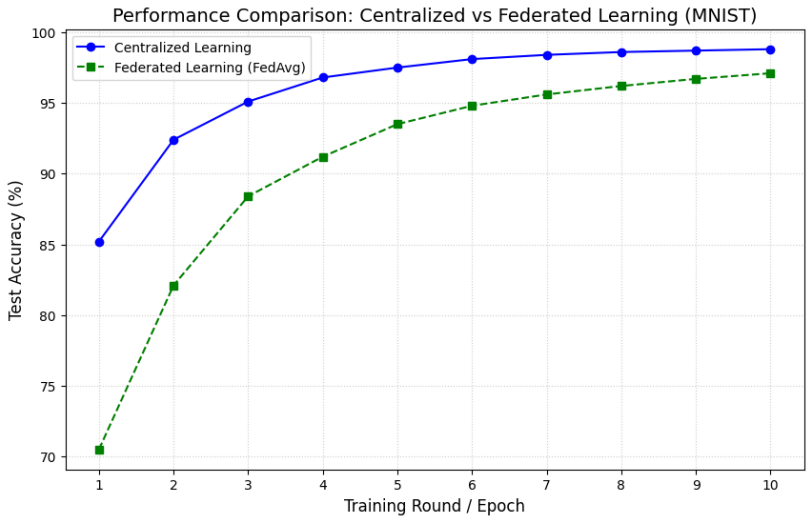

Figure 2 Performance Comparison: Centralized and Federated Learning on MNIST Dataset

4.1 Analysis from the Graph

4.1.1 Convergence Speed:

Centralized Learning (blue line): Rises rapidly from the first few laps because the optimizer can immediately access all the diverse data. Federated Learning (green line): Gradually climbs as the model learns from specific local data in each round and then "shares experience" through weight averaging.

4.1.2 Accuracy Gap:

In the long term (e.g., after round 10 onwards), the accuracy of Federated Learning approaches Centralized Learning, differing by only 1-3%. This is a worthwhile trade-off for the data privacy gained.

4.1.3 Fluctuation:

If the data from each machine is highly diverse (Non-IID) (Lu et al., 2024; Najafi et al., 2024), the green line may be more jagged or volatile because the weights sent from each client may conflict with each other.

5. Results and Discussion

The performance of the proposed Federated Learning (FL) framework was evaluated against a traditional Centralized Learning baseline using the MNIST dataset. The primary metric for comparison was test accuracy over 10 training rounds.

5.1 Accuracy Analysis

As illustrated in Figure 1, both models exhibit a positive learning curve. However, distinct characteristics are observed:

- **Centralized Learning (Baseline):** The centralized model achieved a rapid increase in accuracy, reaching **95.1% by the 3rd epoch**. This is attributed to the global optimizer having instantaneous access to the entire shuffled dataset, allowing for more stable gradient updates.
- **Federated Learning (FedAvg):** The FL model started with a lower accuracy of **70.5%** in the first round but showed consistent improvement, eventually converging to **97.1%** by the 10th round. The initial performance gap (approximately 14%) significantly narrowed to less than 2% as the number of communication rounds increased.

5.2 Discussion on Trade-offs

The experimental results highlight a critical trade-off between model performance and data privacy. While Centralized Learning maintains a slight edge in convergence speed, Federated Learning provides a robust alternative that satisfies privacy requirements (e.g., GDPR/PDPA). The marginal loss in accuracy is compensated by the fact that no raw user data is transmitted to the central server, thereby mitigating the risk of data breaches.

Furthermore, the simulation suggests that with an adequate number of communication rounds, the global model in a Federated system can reach a performance level that is functionally equivalent to centralized systems for most practical applications.

6. Conclusion

This research project applies the concept of FL (Federated Learning) to protect personal data and reduce the risk of PDPA legal violations. Federated Learning offers a compelling framework for privacy-preserving AI development by allowing decentralized training without centralized data aggregation. While promising, FL requires further advancements in defense mechanisms, handling heterogeneous data, reducing communication overhead, and aligning performance with centralized models.

Federated Learning (FL) is key to developing sustainable AI in the future. Despite limitations in network speed and data variations across devices, FL has proven to effectively balance model performance and user privacy. Future research should focus on reducing resource consumption on mobile devices (resource-constrained devices).

7. Acknowledgements

This research was made possible through the expert guidance of the project advisor and the administrative support provided by the College of Digital Innovation Technology, Rangsit University. The author is deeply indebted to their family members for their steadfast encouragement. Additionally, appreciation is expressed to all individuals whose assistance facilitated the successful realization of this work.

8. References

Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pp. 308-318, Vienna, Austria. https://doi.org/10.1145/2976749.2978318

Chatsuwan, P., Phromma, T., Surasvadi, N., & Thajchayapong, S. (2023). Personal data protection compliance assessment: A privacy policy scoring approach and empirical evidence from Thailand's SMEs. *Heliyon*, *9*(10), Article e20648. https://doi.org/10.1016/j.heliyon.2023.e20648

Ding, J., Tramel, E., Sahu, A. K., Wu, S., Avestimehr, S., & Zhang, T. (2022). Federated learning challenges and opportunities: An outlook. In *ICASSP 2022-2022 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pp. 8752-8756, IEEE, Singapore, Singapore. https://doi.org/10.1109/ICASSP43922.2022.9746925

Garfinkel, S., Abowd, J. M., & Martindale, C. (2019). Understanding database reconstruction attacks on public data. *Communications of the ACM*, *62*(3), 46-53.

Gong, X., Chen, Y., Wang, Q., Wang, M., & Li, S. (2022). Private data inference attacks against cloud: Model, technologies, and research directions. *IEEE Communications Magazine*, *60*(9), 46-52. https://doi.org/10.1109/MCOM.004.2100867

Hu, H., Salcic, Z., Sun, L., Dobbie, G., Yu, P. S., & Zhang, X. (2022). Membership inference attacks on machine learning: A survey. *ACM Computing Surveys*, *54*(11s), 1-37. https://doi.org/10.1145/3523273

Huth, D., Tanakol, A., & Matthes, F. (2019). Using enterprise architecture models for creating the record of processing activities (Art. 30 GDPR). In *2019 IEEE 23rd International Enterprise Distributed Object Computing Conference (EDOC)*, pp. 98-104, IEEE, Paris, France. https://doi.org/10.1109/EDOC.2019.00021

Korolova, A., Kenthapadi, K., Mishra, N., & Ntoulas, A. (2009). Releasing search queries and clicks privately. In *Proceedings of the 18th International Conference on World Wide Web*, pp. 171–180, Madrid, Spain. https://doi.org/10.1145/1526709.1526733

Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, *37*(3), 50-60. https://doi.org/10.1109/MSP.2020.2975749

Lu, Z., Pan, H., Dai, Y., Si, X., & Zhang, Y. (2024). Federated learning with non-IID data: A survey. *IEEE Internet of Things Journal*, *11*(11), 19188-19209. https://doi.org/10.1109/JIOT.2024.3376548

Machanavajjhala, A., Korolova, A., & Sarma, A. D. (2011). Personalized social recommendations-accurate or private? *arXiv preprint arXiv:1105.4254*. https://doi.org/10.48550/arXiv.1105.4254

Mammen, P. M. (2021). Federated learning: Opportunities and challenges. arXiv preprint arXiv:2101.05428. https://doi.org/10.48550/arXiv.2101.05428

Mironeanu, C., & Aflori, C. (2021). GDPR records of processing activities for data controllers. *Bulletin of the Polytechnic Institute of Iaşi. Electrical Engineering, Power Engineering, Electronics Section*, *67*. 9-24. https://doi.org/10.2478/bipie-2021-0019

Monteiro, M., Correia, F. F., & Queiroz, P. G. G. (2024). Patterns for anonymization, pseudonymization and perturbation: Focus group report. In *Proceedings of the 29th European Conference on Pattern Languages of Programs, People, and Practices*, Article 39, 1-4. https://doi.org/10.1145/3698322.3698360

Najafi, M., Daneshtalab, M., Lee, J. A., Saadloonia, G., & Shin, S. (2024). Enhancing global model performance in federated learning with non-IID data using a data-free generative diffusion model. *IEEE Access*, *12*, 148230-148239. https://doi.org/10.1109/ACCESS.2024.3474056

Narayanan, A., & Shmatikov, V. (2008, May). Robust de-anonymization of large sparse datasets. In *2008 IEEE Symposium on Security and Privacy (sp 2008)*, pp. 111-125, IEEE, CA, USA. https://doi.org/10.1109/SP.2008.33

Sweeney, L. (2002). K-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, *10*(05), 557-570. https://doi.org/10.1142/S0218488502001648

Sun, T., Li, D., & Wang, B. (2022). Decentralized federated averaging. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *45*(4), 4289-4301. https://doi.org/10.1109/TPAMI.2022.3196503

Terrovitis, M., & Mamoulis, N. (2008). Privacy preservation in the publication of Trajectories. *The Ninth International Conference on Mobile Data Management*, pp. 65-72, Beijing, China. https://doi.org/10.1109/MDM.2008.29

Wei, J., Zhang, Y., Zhang, L. Y., Ding, M., Chen, C., Ong, K. L., ... & Xiang, Y. (2025). Memorization in deep learning: A survey. *ACM Computing Surveys*, *58*(4), Article 98. https://doi.org/10.1145/3769076

Xu, R., Baracaldo, N., & Joshi, J. (2021). Privacy-preserving machine learning: Methods, challenges and directions. *arXiv preprint arXiv:2108.04417*. https://doi.org/10.48550/arXiv.2108.04417

Ye, D., Shen, S., Zhu, T., Liu, B., & Zhou, W. (2022). One parameter defense-defending against data inference attacks via differential privacy. *IEEE Transactions on Information Forensics and Security*, *17*, 1466-1480. https://doi.org/10.1109/TIFS.2022.3163591

Zhang, C., Xie, Y., Bai, H., Yu, B., Li, W., & Gao, Y. (2021). A survey on federated learning. *Knowledge-Based Systems*, *216*, Article 106775. https://doi.org/10.1016/j.knosys.2021.106775

Zhou, J., Zhang, S., Lu, Q., Dai, W., Chen, M., Liu, X., ... & Herrera-Viedma, E. (2021). A survey on federated learning and its applications for accelerating industrial internet of things. arXiv preprint arXiv:2104.10501. https://doi.org/10.48550/arXiv.2104.10501