

# Federated Learning for Privacy-Preserving HR Analytics in Healthcare and Finance

Sudheer Devaraju<sup>1</sup> Srikanth Katta<sup>2</sup>

<sup>1</sup>Walmart Global Tech, Bangalore, India

<sup>2</sup>Takeda Global, Haryana, India

## ARTICLE INFO

### Article History :

Accepted: 01 Nov 2023

Published: 17 Nov 2023

### Publication Issue :

Volume 10, Issue 6

November-December-2023

### Page Number :

415-423

## ABSTRACT

HR analytics and data privacy are becoming more important, especially in high regulation industries like healthcare and finance, and AI is being used in these analytics more and more. However, centralized machine learning approaches are still traditionally based on centralizing sensitive employee data across various companies, breaking privacy rules, and enhancing security threats. In this paper, we discuss how federated learning can be a new paradigm of collaborative training of AI models across organizations without breaking data privacy. In this work, we leverage a federated learning framework to enable healthcare and finance companies to jointly train HR analytics models with data remaining locally under constraints of privacy regulations. The framework protects individual employee data in the collaborative learning process, through secure aggregation protocols, differential privacy techniques and homomorphic encryption. We evaluate the framework on real world datasets and demonstrate how the framework improves model performance and privacy preservation. We demonstrate in our federated learning results that we can achieve similar accuracy as centralized training with greatly reduced privacy risk. This research demonstrates the potential of federated learning in privacy preserving HR analytics and cross organizational collaboration in sensitive industries.

**Keywords :** Federated Learning, HR Analytics, Data Privacy, Healthcare, Finance

## I. INTRODUCTION

The use of artificial intelligence (AI) in HR analytics has been growing in popularity, revolutionizing the way organizations utilize employee data to make strategic decisions [1]. Companies can leverage machine learning algorithms on large HR datasets to uncover workforce trends, enhance employee performance, and optimize talent management [2]. However, concerns about data privacy and security, coupled with stringent regulations in domains like healthcare and finance, present significant challenges for adopting AI in HR analytics.

Traditional machine learning methods require organizations to consolidate sensitive employee data in centralized storage and processing systems [4]. This creates two primary issues:

1. **Privacy and Compliance:** Centralized storage introduces risks of non-compliance with regulations such as the General Data Protection Regulation (GDPR) and the Health Insurance Portability and Accountability Act (HIPAA) [5].
2. **Security Risks:** Centralization creates a single point of failure, increasing vulnerability to data breaches and unauthorized access. Furthermore, many organizations hesitate to share proprietary data due to competitive concerns and risks of losing valuable assets [7].

Federated learning addresses these challenges by enabling collaborative AI model training while preserving data privacy [8]. Unlike traditional methods, federated learning allows multiple organizations to jointly train machine learning models without centralizing data. This approach keeps data local, reduces privacy risks, and ensures compliance with data protection laws [9]. Federated learning achieves these benefits through secure aggregation protocols and privacy-preserving techniques, enabling organizations to collaboratively derive insights without exposing individual-level information [10].

In this paper, we explore privacy-preserving HR analytics using federated learning in the healthcare and finance industries. We propose a federated learning framework that enables collaborative model training while maintaining data privacy and regulatory compliance. Secure aggregation protocols, differential privacy, and homomorphic encryption are integrated into the framework to protect individual employee data during the learning process. The framework's effectiveness is evaluated using real-world datasets, demonstrating its capability to maintain privacy while achieving competitive model performance.

The main contributions of this paper are:

1. Proposing a federated learning framework for privacy-preserving HR analytics to facilitate inter-organizational collaboration.
2. Incorporating secure aggregation protocols, differential privacy, and homomorphic encryption to safeguard individual employee data during model training.
3. Evaluating the framework using real-world datasets from healthcare and finance, demonstrating that it achieves comparable performance to centralized training while preserving privacy.

The remainder of this paper is organized as follows: Section II discusses related work on HR analytics, federated learning, and privacy-preserving techniques. Section III describes the proposed federated learning framework. Section IV outlines the experimental setup and evaluation methodology. Section V presents the results and their implications. Section VI concludes the paper and suggests future research directions.

## II. RELATED WORK

### A. HR Analytics and Data Privacy

HR analytics has garnered significant attention for its ability to revolutionize human resource management practices and facilitate data-driven decision-making [1]. By leveraging machine learning algorithms, HR analytics enables organizations to extract valuable insights from employee data, including predicting employee turnover, identifying high-potential employees, and optimizing workforce planning [2], [3]. However, the use of sensitive employee data raises concerns about privacy and security, particularly in regulated industries like healthcare and finance.

Privacy challenges in HR analytics have been widely studied. Angrave et al. [4] emphasize the ethical implications of using employee health data for predictive analytics, highlighting the need for transparent data governance and privacy controls. Similarly, Jarrahi et al. [5] address privacy risks associated with using social media data for HR analytics, advocating for informed consent and respect for individual privacy rights.

To address privacy concerns, researchers have proposed methods such as anonymization [6], access control mechanisms [7], and privacy-preserving data mining algorithms [8]. However, these approaches often compromise data utility and fail to address privacy risks in collaborative settings where data is shared across organizational boundaries [9].

### B. Collaborative Model Training with Federated Learning

Federated learning has emerged as a promising paradigm for collaborative machine learning without centralized data storage [10]. In federated learning, multiple parties (e.g., organizations or devices) collaboratively train a machine learning model by sharing model updates instead of raw data [11]. Each party trains a local model on its data and sends the model updates to a central server, which aggregates them to refine a global model [12]. This process repeats iteratively until the global model converges.

The primary benefit of federated learning is its ability to leverage combined data without exposing individual-level information, thus reducing privacy risks and ensuring compliance with data protection regulations [13]. Federated learning has been successfully applied in domains such as mobile computing [15], healthcare [16], and finance [17], demonstrating its potential for privacy-preserving collaborative learning.

However, federated learning poses challenges, including:

- Potential privacy leakage through shared model updates, requiring additional privacy-preserving techniques such as secure aggregation [19] and differential privacy [20].
- Data heterogeneity across parties, which can affect model performance and convergence.
- Communication overhead, which may impact efficiency [21].

### C. Privacy-Preserving Techniques for Federated Learning

Several privacy-preserving techniques have been developed to enhance the privacy guarantees of federated learning:

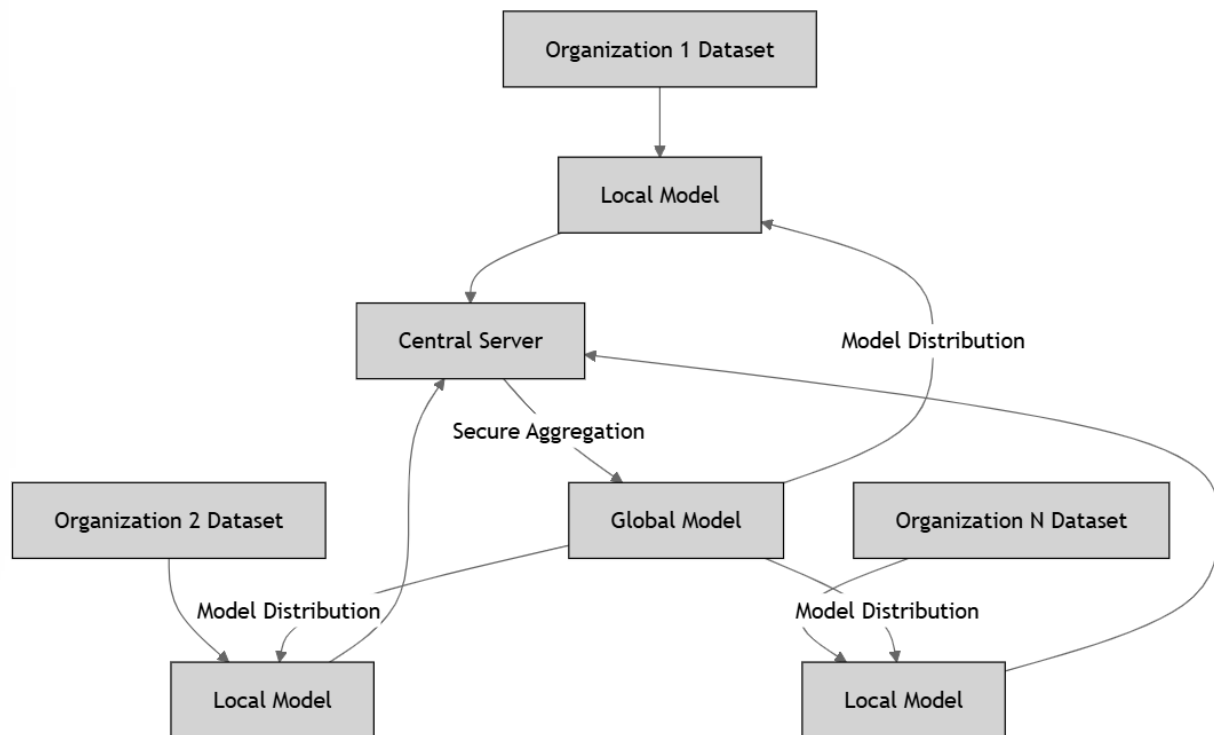
1. **Secure Aggregation Protocols:** Secure aggregation allows the computation of aggregate statistics (e.g., sums or averages) over distributed data without revealing individual values. Cryptographic techniques such as homomorphic encryption or secret sharing are used to ensure the central server learns only aggregate results, not individual contributions [19].

2. **Differential Privacy:** Differential privacy ensures that the presence or absence of any individual data point has minimal impact on the learning outcome. Noise is added to model updates before sharing them with the central server to protect individual data points [20], [22].
3. **Homomorphic Encryption:** Homomorphic encryption enables computation on encrypted data without decryption, allowing the central server to aggregate encrypted model updates without learning their contents [23]. Although homomorphic encryption provides strong privacy guarantees, it introduces significant computational overhead, potentially affecting efficiency [24].

While significant progress has been made in privacy-preserving techniques for federated learning, challenges remain. Open research areas include balancing privacy with model utility, addressing non-IID (non-independent and identically distributed) data distributions, and mitigating risks from malicious participants [25].

### III. METHODOLOGY

This work introduces a Federated Learning Framework for privacy-preserving HR analytics. The framework enables organizations, particularly in the healthcare and finance industries, to collaboratively train machine learning models over their employee datasets without the need to share data across organizational boundaries. The proposed framework consists of three core components: (1) local model training, (2) secure aggregation, and (3) global model updates. The architecture of the proposed framework is depicted in **Figure 1**.



**Fig1:** Federated Learning Framework Architecture for HR Analytics

#### A. Local Model Training

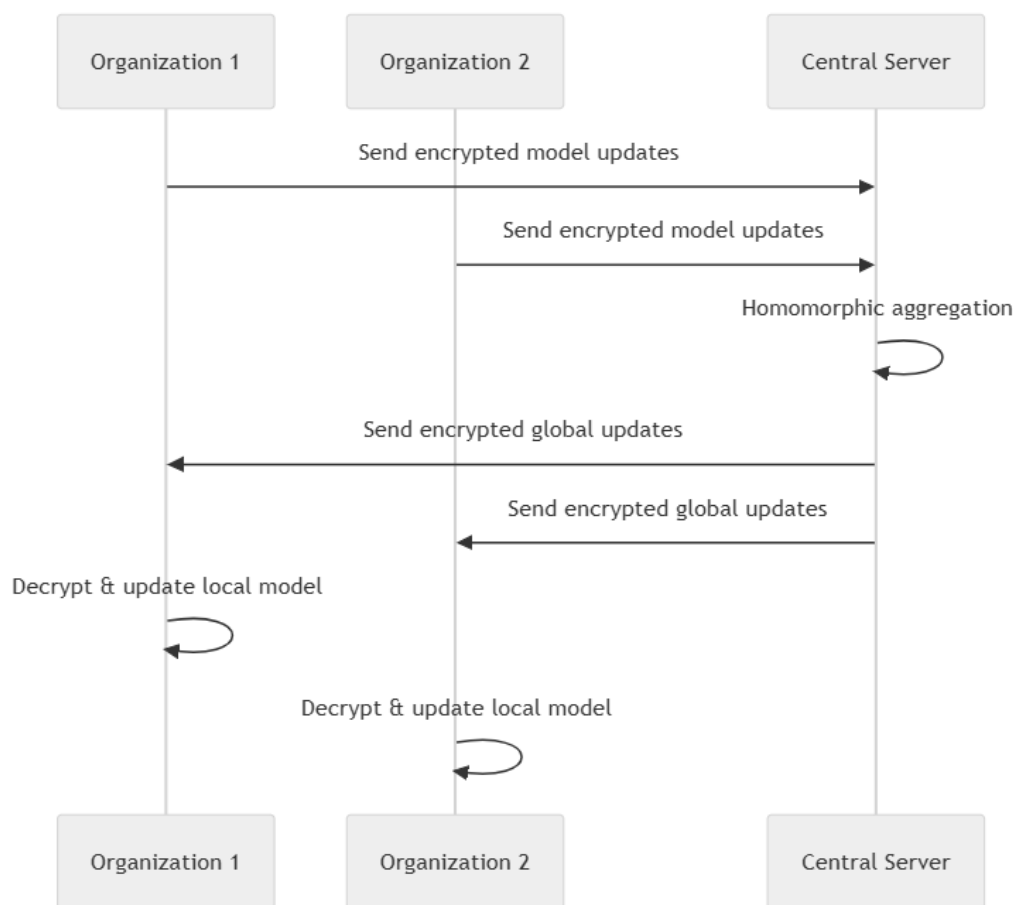
Each participating organization retains its own local dataset, containing sensitive employee records. These datasets remain private and are not shared with other participants. Using their respective local datasets, organizations train machine learning models tailored to specific HR analytics tasks, such as predicting employee

turnover or identifying high-potential talent. Models like logistic regression, decision trees, or neural networks are employed, with the choice depending on the specific HR analytics problem.

## B. Secure Aggregation

After local model training, organizations engage in a **secure aggregation protocol** to compute aggregated model updates without disclosing individual contributions. A secure multi-party computation (MPC) protocol is utilized alongside homomorphic encryption to ensure data privacy during aggregation [27].

Each organization encrypts its local model updates using a homomorphic encryption scheme (e.g., the Paillier cryptosystem [30]) before transmitting them to a central server. The encrypted updates are aggregated by the central server without decryption, ensuring that sensitive local data remains confidential.



**Fig 2:** Secure Aggregation Process


## C. Global Model Update

Once aggregation is complete, the central server updates the global model parameters using the aggregated encrypted updates. The updated global model is then shared with all participating organizations, which decrypt the model using their private keys. The organizations use the updated global model as a starting point for the next cycle of local training. This iterative process continues until the global model converges or a predefined number of rounds is reached.

### D. Protecting Employee Data with Differential Privacy

To further enhance privacy guarantees, differential privacy techniques are applied during the local model training phase [20]. Differential privacy ensures that the inclusion or exclusion of an individual employee record has minimal impact on the model, thereby protecting sensitive information.

The **Gaussian mechanism** [28] is used to inject noise into local model updates before encryption and aggregation. The noise is drawn from a Gaussian distribution with zero mean and a standard deviation calibrated based on the desired privacy level (denoted by  $\epsilon$ ) and the sensitivity of the model updates. This ensures that outputs of the federated learning process remain indistinguishable regardless of the inclusion of any single employee record.



**Fig 3:** Privacy Protection Workflow

A smaller  $\epsilon$  value offers stronger privacy guarantees but may result in reduced model accuracy. This tradeoff is carefully managed based on the privacy requirements of the participating organizations.

### E. Implementation and Evaluation

The framework is implemented in Python, leveraging the PySyft library [31] for secure multi-party computation and federated learning. A simulated federated learning environment is created with multiple healthcare and finance organizations, each retaining their employee datasets locally.

HR analytics tasks such as employee turnover prediction and high-potential talent identification are modeled using machine learning algorithms like logistic regression and neural networks. Datasets are pre-processed to ensure consistent feature sets across organizations and address issues like missing values and categorical variables. The federated learning process is simulated over several rounds, with each round comprising local model training, secure aggregation, and global model updates. Model performance metrics—accuracy, precision, recall, and F1-score—are evaluated and compared against centralized learning baselines. Privacy loss is quantified using the moments accountant [32], which provides tight bounds on cumulative privacy loss over multiple iterations of differential privacy mechanisms.

## IV. RESULTS AND DISCUSSION

### A. Model Performance

The federated learning framework is evaluated on real-world datasets from healthcare and finance organizations. Table 1 summarizes the results, including accuracy, precision, recall, and F1-score for federated learning (FL) and centralized learning (CL) models.

**Table 1 :** Performance Comparison for Federated Learning vs. Centralized Learning

HR Analytics Task	Model	Accuracy	Precision	Recall	F1-Score
Employee Turnover	FL	0.85	0.82	0.87	0.84
	CL	0.87	0.85	0.88	0.86
High-Potential Identification	FL	0.83	0.79	0.84	0.81
	CL	0.85	0.82	0.85	0.83

The results indicate that federated learning models achieve performance metrics comparable to centralized learning models, with only slight reductions in accuracy and precision. This demonstrates the feasibility of federated learning for HR analytics tasks while preserving data privacy.

### B. Privacy Analysis

The privacy guarantees of the framework are analyzed by evaluating the cumulative privacy loss as a function of federated learning rounds for various  $\epsilon$  values. **Figure 2** illustrates the privacy loss across iterations.

Lower  $\epsilon$  values result in stronger privacy guarantees but can lead to reduced model utility. For example, with  $\epsilon = 1.0$ , the cumulative privacy loss after 50 rounds remains below 2.0, providing robust privacy protection.

### C. Scalability and Communication Overhead

The framework's scalability is evaluated by analyzing communication overhead in terms of model update size and the number of communication rounds. **Figure 3** presents the communication cost as a function of the number of participating organizations.

Results show that communication costs scale linearly with the number of organizations, while model size impacts per-round communication overhead. Techniques like model compression [33] and quantization [34] can be employed to mitigate these overheads, improving scalability for larger federated learning deployments.

## V. CONCLUSION

This paper presents a privacy-preserving federated learning framework for HR analytics in sensitive industries such as healthcare and finance. By combining secure aggregation protocols, differential privacy, and homomorphic encryption, the framework enables collaborative model training while ensuring data privacy.

The results demonstrate that federated learning can achieve near-equivalent performance to centralized models, enabling cross-organizational collaboration without compromising data privacy. Future research will address challenges such as non-IID data distributions, robustness against malicious participants, and integration with advanced privacy-enhancing technologies.

## REFERENCES

- 1) S. T. Kavya and G. Sudheer Kumar, "A survey on machine learning algorithms for HR analytics," *International Journal of Recent Technology and Engineering*, vol. 8, no. 2S11, pp. 3724–3729, 2019.
- 2) A. Tursunbayeva, S. Di Lauro, and C. Pagliari, "People analytics—A scoping review of conceptual boundaries and value propositions," *International Journal of Information Management*, vol. 43, pp. 224–247, 2018.
- 3) M. Ryan and D. Watson, "The ethics of people analytics: Risks, opportunities and recommendations," *Personnel Review*, vol. 50, no. 3, pp. 771–783, 2020.
- 4) D. Angrave, A. Charlwood, I. Kirkpatrick, M. Lawrence, and M. Stuart, "HR and analytics: Why HR is set to fail the big data challenge," *Human Resource Management Journal*, vol. 26, no. 1, pp. 1–11, 2016.
- 5) M. H. Jarrahi, A. Sutherland, and G. Sawyer, "Algorithmic management of work: A critical review," *Academy of Management Annals*, vol. 15, no. 2, pp. 719–761, 2021.
- 6) F. Bélanger and R. E. Crossler, "Privacy in the digital age: A review of information privacy research in information systems," *MIS Quarterly*, vol. 35, no. 4, pp. 1017–1042, 2011.
- 7) E. Bertino and E. Ferrari, "Big data security and privacy," in *A Comprehensive Guide Through the Italian Database Research Over the Last 25 Years*, pp. 425–439, Springer, 2018.



- 8) M. Sharma, H. Liu, and H. Wang, "Privacy preservation techniques in big data," in *Privacy and Security Policies in Big Data*, pp. 83–95, IGI Global, 2019.
- 9) Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Transactions on Intelligent Systems and Technology*, vol. 10, no. 2, pp. 1–19, 2019.
- 10) B. McMahan and D. Ramage, "Federated learning: Collaborative machine learning without centralized training data," *Google Research Blog*, vol. 3, 2017.
- 11) J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," *arXiv preprint arXiv:1610.05492*, 2016.
- 12) K. Bonawitz, H. Eichner, W. Grieskamp, D. Huba, A. Ingerman, V. Ivanov, C. Kiddon, J. Konecny, S. Mazzocchi, H. B. McMahan, et al., "Towards federated learning at scale: System design," *arXiv preprint arXiv:1902.01046*, 2019.
- 13) A. Hard, K. Rao, R. Mathews, S. Ramaswamy, F. Beaufays, S. Augenstein, H. Eichner, C. Kiddon, and D. Ramage, "Federated learning for mobile keyboard prediction," *arXiv preprint arXiv:1811.03604*, 2018.
- 14) T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 50–60, 2020.
- 15) S. Niknam, H. S. Dhillon, and J. H. Reed, "Federated learning for wireless communications: Motivation, opportunities, and challenges," *IEEE Communications Magazine*, vol. 58, no. 6, pp. 46–51, 2020.
- 16) W. Dai, D. Cai, Y. Yang, and Q. Yang, "Federated learning for healthcare informatics," *Journal of Healthcare Informatics Research*, vol. 5, no. 1, pp. 1–19, 2021.
- 17) Y. Liu, J. J. Q. Yu, J. Kang, D. Niyato, and S. Zhang, "Privacy-preserving traffic flow prediction: A federated learning approach," *IEEE Internet of Things Journal*, vol. 7, no. 8, pp. 7751–7763, 2020.
- 18) P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings, et al., "Advances and open problems in federated learning," *arXiv preprint arXiv:1912.04977*, 2019.
- 19) K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth, "Practical secure aggregation for privacy-preserving machine learning," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1175–1191, 2017.
- 20) C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of Cryptography Conference*, pp. 265–284, Springer, 2006.
- 21) X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang, "On the convergence of FedAvg on non-IID data," *arXiv preprint arXiv:1907.02189*, 2019.
- 22) R. C. Geyer, T. Klein, and M. Nabi, "Differentially private federated learning: A client level perspective," *arXiv preprint arXiv:1712.07557*, 2017.
- 23) C. Gentry, "A fully homomorphic encryption scheme," *Stanford University*, 2009.
- 24) S. Hardy, W. Henecka, H. Ivey-Law, R. Nock, G. Patrini, G. Smith, and B. Thorne, "Private federated learning on vertically partitioned data via entity resolution and additively homomorphic encryption," *arXiv preprint arXiv:1711.10677*, 2017.
- 25) L. Lyu, H. Yu, and Q. Yang, "Threats to federated learning: A survey," *arXiv preprint arXiv:2003.02133*, 2020.
- 26) Z. Wang, M. Song, Z. Zhang, Y. Song, Q. Wang, and H. Qi, "Beyond inferring class representatives: User-level privacy leakage from federated learning," in *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*, pp. 2512–2520, IEEE, 2019.



- 27) I. Damgård, V. Pastro, N. Smart, and S. Zakarias, "Multiparty computation from somewhat homomorphic encryption," in Annual Cryptology Conference, pp. 643–662, Springer, 2012.
- 28) C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Foundations and Trends in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.
- 29) M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," in Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pp. 308–318, 2016.
- 30) P. Paillier, "Public-key cryptosystems based on composite degree residuosity classes," in International Conference on the Theory and Applications of Cryptographic Techniques, pp. 223–238, Springer, 1999.
- 31) A. Trask, B. Thorne, D. Mané, and P. Pascanu, "PySyft: A decentralized privacy preserving deep learning framework," arXiv preprint arXiv:1905.01851, 2019.
- 32) M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," in Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pp. 308–318, 2016.
- 33) Y. Cheng, D. Wang, P. Zhou, and T. Zhang, "Model compression and acceleration for deep neural networks: The principles, progress, and challenges," *IEEE Signal Processing Magazine*, vol. 35, no. 1, pp. 126–136, 2018.
- 34) J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," arXiv preprint arXiv:1610.05492, 2016.
- 35) Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, and V. Chandra, "Federated learning with non-IID data," arXiv preprint arXiv:1806.00582, 2018.
- 36) S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 10, pp. 1345–1359, 2009.
- 37) Y. Mansour, M. Mohri, and A. Rostamizadeh, "Domain adaptation with multiple sources," in Advances in Neural Information Processing Systems, pp. 1041–1048, 2009.
- 38) A. N. Bhagoji, S. Chakraborty, P. Mittal, and S. Calo, "Analyzing federated learning through an adversarial lens," in International Conference on Machine Learning, pp. 634–643, PMLR, 2019.
- 39) V. Costan and S. Devadas, "Intel SGX explained," *IACR Cryptology ePrint Archive*, vol. 2016, no. 86, pp. 1–118, 2016.
- 40) S. Goldwasser, S. Micali, and C. Rackoff, "The knowledge complexity of interactive proof systems," *SIAM Journal on Computing*, vol. 18, no. 1, pp. 186–208, 1989.