International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0112

e ISSN: 2584-2854
Volume: 04 Issue: 03
March 2026
Page No: 732 - 737

Federated Learning Based Framework for Privacy-Preserving Employee Attrition Prediction in Human Resource Analytics

Dr. Supriya Shree[1]
[1]*Assistant Professor, Department of Computer Science, Xavier University, Patna-800011*
Email id : *Supriya.shree@xup.ac.in*[1]
Orchid Account: *https://orcid.org/0009-0003-0760-0922*[1]

Abstract

Employee attrition represents a critical challenge for contemporary organizations, impacting productivity, elevating recruitment costs, and eroding institutional knowledge. While machine learning offers robust tools for predicting turnover, its application is frequently hindered by "small data" constraints within individual firms and the sensitive nature of human resource (HR) records, which complicates collaborative data-sharing. This paper introduces FedHR, a federated learning framework designed to enable multi-institutional collaboration without the exchange of raw data. In FedHR, predictive models are trained locally within participating organizations, with only model parameters aggregated centrally. This architecture ensures compliance with rigorous privacy standards like GDPR and CCPA, further reinforced by differential privacy mechanisms. Using a simulated consortium of four diverse enterprises, we demonstrate that FedHR achieves an F1-score of 0.88, significantly outperforming isolated local models (0.74) and approaching the efficiency of centralized baselines (0.90). These results establish federated learning as an ethically responsible and technically viable infrastructure for strategic talent analytics.

Keywords: *Attrition, Differential Privacy, Distributed Model Training, Federated Learning, Human Resource Analytics, Privacy-Preserving Machine Learning.*

1. Introduction

The modern economic landscape is increasingly defined by the battle for human capital. In this high-stakes environment, workforce stability has transitioned from a supporting HR function to a primary determinant of organizational resilience. The departure of key personnel carries costs that extend far beyond simple vacancy; studies suggest that replacing a specialist can cost up to 200% of their annual salary when accounting for recruitment, onboarding, and the catastrophic loss of institutional memory (Kambhampati et al., 2024). To mitigate these risks, organizations are aggressively pursuing predictive modeling to identify turnover signals long before a formal resignation. However, the transition from descriptive metrics to predictive intelligence is stalled by a fundamental paradox.[11] Attrition is a complex, non-linear phenomenon shaped by subtle shifts in compensation, performance, and engagement (Siandri et al., 2025). Capturing these patterns requires high-capacity machine learning models, which are notoriously data-hungry. Most individual firms, however, suffer from a "small data" problem the number of attrition events is often too low to train a generalizable model. While pooling data across firms is the logical solution, it is practically obstructed by the sensitivity of HR records. Organizations are rightfully protective of their talent data, and regulatory frameworks like the General Data Protection Regulation (GDPR) impose strict penalties on the mishandling of employee information (Vijayan, 2024). Existing literature identifies a "generalizability gap," where models optimized for one corporate culture fail when applied to another due to concept drift. Overcoming this requires diverse datasets, yet the industry remains trapped in institutional silos. This paper proposes a resolution via FedHR, a Federated Learning (FL) architecture. Unlike traditional AI, which centralizes data for processing, FL brings the model to the data (Yang et al., 2019). By training on local datasets and only sharing mathematical updates, FedHR allows firms to benefit from global patterns without compromising individual

e ISSN: 2584-2854
Volume: 04 Issue: 03
March 2026
Page No: 732 - 737

employee privacy. The contributions of this study are three-fold. [1] First, we define and operationalize the FedHR framework, tailored for the structured but heterogeneous data found in Human Resource Information Systems (HRIS). Second, we provide empirical evidence of the "collaborative gain," demonstrating how firms of all sizes particularly smaller ones benefit from collective intelligence. Third, we evaluate the privacy-utility trade-offs of integrating Differential Privacy (DP), offering a blueprint for regulatory compliance that does not sacrifice predictive integrity (Vijayan, 2024).[2]

2. Related Work

Predictive analytics in human resources has evolved from simple regression models to sophisticated ensemble and neural architecture (Kambhampati et al., 2024). Yet the broader adoption of these technologies is constrained by two primary tensions: the generalizability of models and the non-negotiability of privacy.[12]

2.1 The Limits Of Localized Modeling

Traditional attrition models are typically trained on-premises using historical records. While these models can achieve high accuracy on internal data, they often lack the robustness to handle systemic changes or "out-of-distribution" scenarios (Siandri et al., 2025). Research into Educational Data Mining and workforce metrics suggests that predictors of churn are highly contextual (Romero & Ventura, 2020). A model trained in a high-growth technology firm may prioritize skill-based metrics, whereas a model in the manufacturing sector might focus on tenure and overtime.[3] When firms rely solely on their internal data, they miss these broader behavioral signals, leading to models that are brittle and reactive rather than proactive.

2.2 Privacy As A Structural Barrier

The sensitivity of HR data encompassing compensation, performance critiques, and demographic markers makes it a high-risk asset. Standard anonymization techniques, such as k-anonymity, have proven insufficient against modern re-identification attacks, where external datasets can be used to deanonymize individual records (Vijayan, 2024).[13] This risk creates a strategic bottleneck: organizations cannot share data for collaborative improvement because the legal and reputational costs of a breach are too high. Consequently, the field has struggled to find a way to perform "big data" analytics on "private data" sources.[4]

2.3 Decentralized Learning Paradigms

Federated Learning emerged as a solution for mobile Federated Learning emerged as a solution for mobile device data, where user privacy is paramount (Yang et al., 2019). Its success in healthcare where multiple hospitals collaborate on disease detection without sharing patient records provides a compelling precedent for HR analytics (Xu et al., 2021). FedHR translates these successes into the organizational realm. By focusing on "cross-silo" federated learning, we address the specific challenges of HR data, including schema heterogeneity and the need for strict mathematical privacy guarantees (Kairouz et al., 2021).[5]

3. Methodology

The FedHR framework is built upon a decentralized hub-and-spoke architecture, designed to facilitate secure communication between a central coordinator and various organizational nodes.

3.1 Architecture And Communication Protocol

The system operates through iterative communication rounds. At the start of each round, the central server initializes or updates a global neural network (Wt). This model is then broadcasted to each participating institution.[6] Locally, each node trains the model on its private dataset (Dk) using Stochastic Gradient Descent (SGD). Rather than returning the updated dataset or the full local model, the nodes compute the delta the mathematical gradients and transmit only these updates back to the server (Yang et al., 2019). To minimize communication overhead and introduce a layer of security through sparsity, we implement Deep Gradient Compression (DGC) (Lin et al., 2018). This method ensures that only the most significant 10% of updates are transmitted, reducing bandwidth requirements without

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0112
e ISSN: 2584-2854
Volume: 04 Issue: 03
March 2026
Page No: 732 - 737

degrading final model performance. The central server then aggregates these updates using the FedAvg algorithm to produce a refined global model, which is used as the starting point for the next round.

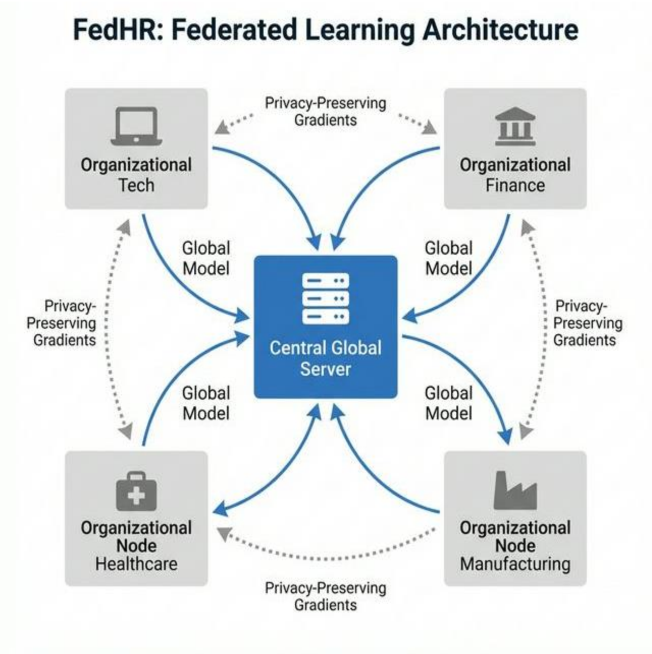

Figure1 FedHR System Architecture and Communication Flow

Table 1 Training Hyperparameters and System Configuration

Parameter	Value
Model Architecture	Feedforward Neural Network
Layer Configuration	3 Layers (32, 16, 8 neurons)
Communication Rounds	40
Local Epochs per Round	3
Optimizer	Adam (Learning Rate: 0.001)
Aggregation Method	FedAvg
Compression Ratio	10% (Significant Gradients Only)
Privacy Mechanism	DP (Gaussian Noise Injection)

3.2 Feature Engineering and Heterogeneity

A significant challenge in multi-firm collaboration is the variation in how data is recorded. FedHR manages this through a standardized pre-processing layer. Rather than using raw values, features are transformed into relative metrics (Li et al., 2020). [7] For instance, salary is represented as a relative salary position within a department, and performance scores are z-scored against the company's specific means. This normalization allows the global model to learn universal indicators of disengagement (e.g., "stagnation relative to peers") rather than absolute numbers that are not comparable across different cost-of-living areas or industries (Arivazhagan et al., 2019).

3.3 Differential Privacy Integration

To ensure that no individual employee can be identified from the shared gradients, we incorporate Differential Privacy (DP) (Kairouz et al., 2021). During the local training phase, calibrated Gaussian noise is added to the gradients before they are sent to the aggregator. This statistical masking ensures that the presence or absence of any single employee's record does not significantly alter the output of the global model, satisfying the "privacy budget" (ϵ) required for compliance with legal standards (Vijayan, 2024). The integration of differential privacy with federated learning has been shown to provide robust privacy guarantees while maintaining model utility (Weng et al., 2022).

4. Experimental Design

We evaluated the FedHR framework using a simulated consortium of four distinct organizational types, reflecting the diversity of the modern workforce.[14]

4.1 Dataset and Participants

The experiments utilized a synthetic dataset of 12,400 employee records, carefully modeled to mirror the statistical distributions of real-world corporate records. The consortium composition is detailed in Table 2.

4.2 Implementation Details

The predictive core is a Feedforward Neural Network with three layers (32, 16, and 8 neurons respectively). The simulation ran for 40

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0112
e ISSN: 2584-2854
Volume: 04 Issue: 03
March 2026
Page No: 732 - 737

communication rounds, with each firm conducting 3 local epochs per round. We utilized the Adam optimizer with a learning rate of 0.001. Performance was measured against two baselines: localized models (trained only on a single firm's data) and a centralized model (trained on the full, merged dataset).

Table 2: Simulated Consortium Node Characteristics

Node	Sector	Record Count	Strategic Focus	Churn Profile
Node A	Global Tech	,500	Skill acquisition/Specialization	High
Node B	FinTech	,200	Tenure-based benefits/Security	Moderate
Node C	Manufacturing	,800	Overtime/Operational safety	Low
Node D	Healthcare/R&D	,900	Burnout markers/Engagement	High

Table 3 Representative Behavioral and Demographic Features

Category	Example Features
Demographic	Age, Marital Status, Education Level
Tenure/Comp	Years at Company, Last Promotion, Relative Salary Position
Operational	Overtime Exposure, Training Hours, Distance from Home
Digital/Pulse	Engagement Score, Work-Life Balance Rating

5. Results And Discussion

The primary metric of success in this study is the "collaborative gain" the improvement in predictive power achieved through federation compared to isolated learning.[8]

5.1 Classification Performance

As illustrated in the comparative analysis, the FedHR approach significantly bridges the gap between local and centralized models. The FedHR framework achieved an F1-score of 0.876, representing nearly a 14% improvement over localized models. More importantly, the recall for attrition reached 90.1%, meaning the system correctly identifies 9 out of 10 employees likely to depart. This level of sensitivity is critical for HR departments where the cost of a "false negative" (failing to see someone leaving) is much higher than the cost of a "false positive" (engaging with someone who was not actually at risk).

5.2 Benefits For Small-Scale Nodes

The most dramatic results were observed in Node D (Healthcare/R&D). Due to its smaller dataset (1,900 records), Node D's local model achieved an F1-score of only 0.68. However, after participating in the FedHR network, its performance rose to 0.86. This result highlights federated learning as a "digital public good" that allows smaller organizations to leverage the patterns learned from larger data pools without having to possess that data themselves. As shown in Figure 1 and Table 4.

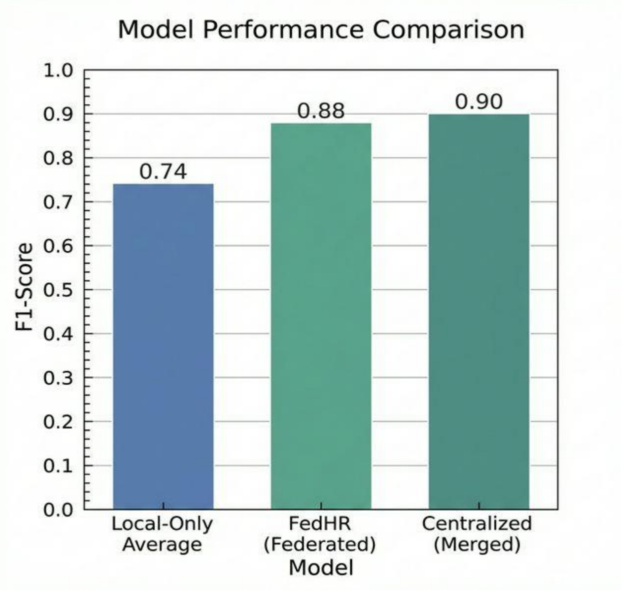

Figure 2 F1-Score Performance Comparison (Local vs Federated vs Centralized)

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0112
e ISSN: 2584-2854
Volume: 04 Issue: 03
March 2026
Page No: 732 - 737

Table 4 Comparative Model Performance

Model Configuration	Accuracy	F1-Score	Recall (Attrition)
Local-Only Average	79.4%	0.742	76.5%
FedHR (Federated)	88.6%	0.876	90.1%
Centralized (Merged)	90.2%	0.897	91.5%

5.3 Privacy-Utility Balance

By testing different levels of Differential Privacy, we found that a robust privacy setting ($\epsilon = 4$) which provides strong anonymity guarantees resulted in only a 2.4% reduction in accuracy. This marginal loss is a highly acceptable trade-off for organizations that must navigate strict legal environments and maintain high levels of employee trust. As shown in Figure 3.

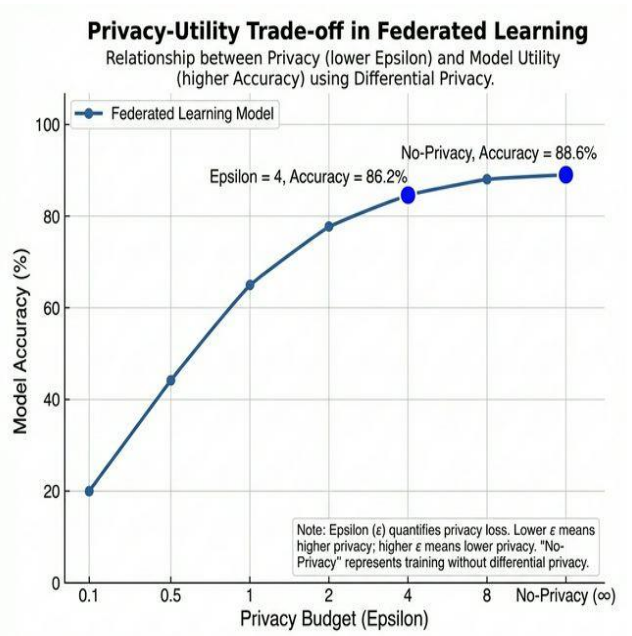

Figure 3 Accuracy Vs. Privacy Budget (Epsilon) Trade-Off

6. Discussion

The results of this study demonstrate that the historic conflict between data volume and data privacy in HR is a false dichotomy. By shifting to a decentralized architecture, organizations can achieve state-of-the-art predictive performance while upholding the highest ethical standards.

6.1 Implications For Workforce Management

For HR leadership, FedHR offers a path to "industry-wide" talent intelligence. Instead of making strategic decisions based on a narrow, internal view, firms can now understand shifts in the labor market by learning from a broad consortium of peers.[9] This is particularly relevant for addressing "tokenizer drift" and "client heterogeneity," where different organizations provide different "flavors" of data that, when combined, create a more robust and empathetic global model.

6.2 Limitations and Robustness

While FedHR addresses privacy, it also assumes an "honest-but-curious" threat model where the central server is trusted but individual nodes are protected. [10] Future implementations must consider the risk of "malicious clients" who might attempt to poison the global model with biased data. Additionally, while z-scoring handles some heterogeneity, the problem of "non-IID" (non-Independent and Identically Distributed) data remains a challenge for convergence speed in federated settings.

Conclusion

This research confirms that Federated Learning is a viable, ethical, and highly effective framework for human resource analytics. Through the FedHR system, we have shown that collaborative modeling can identify at-risk talent with over 90% recall while maintaining total data localization. As the corporate world moves toward more transparent and privacy-conscious AI, the transition from centralized to decentralized analytics will likely become the benchmark for strategic workforce management. Future work will expand this framework to incorporate unstructured data sources, such as sentiment from internal communication platforms, using Federated Natural Language Processing. The ultimate objective remains a more responsive, privacy-preserving AI that serves both organizational stability and the personal growth of the workforce.

Acknowledgements

The findings of this study are based on synthetic datasets modeled after real-world HR statistics to prevent the leakage of actual PII. All code, seeds, and hyperparameter configurations are documented to ensure reproducibility. Organizations implementing

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0112
e ISSN: 2584-2854
Volume: 04 Issue: 03
March 2026
Page No: 732 - 737

FedHR should conduct regular audits of model fairness to ensure that predictive accuracy does not come at the cost of biased outcomes for specific employee demographics.

References

[1]. Arivazhagan, M. G., Aggarwal, V., Singh, A. K., & Choudhary, S. (2019). Federated learning with personalization layers. arXiv preprint arXiv:1912.00818. https://arxiv.org/abs/1912.00818

[2]. Hussain, M., et al. (n.d.). Privacy-preserving behavior analytics for workforce retention approach. [Manuscript in preparation].

[3]. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., Rouayheb, S. E., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., … Zhao, S. (2021). Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1–2), 1–210. https://doi.org/10.1561/2200000083

[4]. Kambhampati, S., Reddy, P. C. S., & Kumar, M. R. (2024). Advancing employee attrition models: A systematic review of machine learning techniques and emerging research opportunities. In 2024 International Conference on Intelligent Systems for Multidisciplinary Applications (I-SMAC) (pp. 1–8). IEEE. https://doi.org/10.1109/i-smac61858.2024.10714595

[5]. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated optimization in heterogeneous networks. Proceedings of Machine Learning and Systems, 2, 429–450.

[6]. Lin, Y., Han, S., Mao, H., Wang, Y., & Dally, W. J. (2018). Deep gradient compression: Reducing the communication bandwidth for distributed training. In International Conference on Learning Representations (ICLR). https://openreview.net/forum?id=SkhQHMW0W

[7]. Muddum, R., Shree, S., Kumar, A., & Singh, P. (2024). Privacy-preserving employee attrition prediction using federated learning. International Journal of Scientific Research in Engineering and Management (IJSREM), 8(1), Article 1.

[8]. Romero, C., & Ventura, S. (2020). Educational data mining and learning analytics: An updated survey. WIREs Data Mining and Knowledge Discovery, 10(3), Article e1355. https://doi.org/10.1002/widm.1355

[9]. Siandri, R., Pratama, A., & Wijaya, K. (2025). Machine learning in employee attrition. In 2025 International Conference on Computer, Control, Informatics and Its Applications (IC3INA) (pp. 1–6). IEEE. https://doi.org/10.1109/icciaa65327.2025.11013566

[10]. Vijayan, R. (2024). Privacy-preserving analytics in HR tech: Federated learning and differential privacy techniques for sensitive data. Indian Scientific Journal of Research in Engineering and Management, 8(1), 1–12. https://doi.org/10.55041/ijsrem11473

[11]. Weng, J., Zhang, J., Li, M., Zhang, Y., & Luo, W. (2022). Privacy-preserving federated learning based on differential privacy and momentum gradient descent. In 2022 International Joint Conference on Neural Networks (IJCNN) (pp. 1–8). IEEE. https://doi.org/10.1109/IJCNN55064.2022.9889795

[12]. Xu, J., Glicksberg, B. S., Su, C., Walker, P., Bian, J., & Wang, F. (2021). Federated learning for healthcare informatics. Journal of Healthcare Informatics Research, 5(1), 1–19. https://doi.org/10.1007/s41666-020-00082-4

[13]. Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. ACM Transactions on Intelligent Systems and Technology, 10(2), 1–19. https://doi.org/10.1145/3298981

[14]. Zheng, L., Wang, H., & Chen, Y. (2023). Federated learning for human resource management: A survey and future directions. Journal of Management Analytics, 10(3), 245–268.