RESEARCH PAPER

Privacy-Preserving Federated Deep Learning for Multi-Class Obesity Risk Stratification in Tamil Nadu Territorial Region.

Rajeswari R[1], Janarthanam S[2]*, Shanthakumar M[3]

[1]Ph.D Researcher, Department of Computer Science, Kamban College of Arts and Science, Coimbatore.
[2]*Associate Professor, Department of Computer Applications, School of Science and Computer Studies, CMR University, Bengaluru.
[3]Assistant Professor, Department of Computer Science, Kamban College of Arts and Science, Coimbatore.

Corresponding author:
Janarthanam S

ABSTRACT
The Obesity prevalence in Tamil Nadu has escalated to 41.3% in males and 27.3% in females, representing a significant public health burden. Centralised data mining approaches for obesity prediction encounter substantial obstacles, including regulatory constraints, institutional data silos, and patient privacy concerns. This investigation presents a privacy-preserving federated deep learning framework for stratifying obesity risk across distributed hospital networks in Tamil Nadu using electronic health records, advancing data-driven public health surveillance while maintaining data sovereignty. A federated learning architecture was implemented across simulated multi- institutional healthcare networks representing diverse urban, semi-urban, and rural settings. The framework integrated data from 724,115 individuals derived from national health surveys, combined with synthesised clinical features. The architecture employed deep neural networks augmented with differential privacy mechanisms and cryptographically secure aggregation protocols. Obesity risk was stratified into four categories through pattern recognition and predictive modelling approaches inherent to advanced data mining. The federated model demonstrated superior performance metrics with 94.2% accuracy and AUC-ROC of 0.963, substantially exceeding centralised approaches (92.3%, AUC 0.948) and conventional federated averaging (91.8%, AUC 0.941). Model convergence occurred within 32 communication rounds compared to 45 rounds required for standard federated averaging. Feature importance analysis identified body mass index (18.5%), waist circumference (17.2%), and chronological age (14.8%) as principal predictive factors. Population-level risk stratification revealed 14.5% classified as very high risk and 24.7% as high risk, facilitating targeted clinical intervention strategies. Federated deep learning architecture successfully enables population-level obesity risk prediction while simultaneously preserving patient privacy and maintaining institutional data autonomy. This approach demonstrates considerable scalability potential for distributed health surveillance and evidence-based intervention programming across Tamil Nadu's decentralised healthcare infrastructure.

Keywords: Federated learning, Obesity prediction, Electronic health records, Deep learning, Risk stratification, Data mining, Differential privacy, Healthcare networks
How to cite this article: R R, Janarthanam S, Shanthakumar M, Privacy-Preserving Federated Deep Learning for Multi-Class Obesity Risk Stratification in Tamil Nadu Territorial Region..Int J Drug Deliv Technol. 2026; 16(8s): 360-372; DOI: 10.25258/ijddt.16.8s.48
Source of support: Nil.
Conflict of interest: None

INTRODUCTION

1.1 Epidemiological Significance of Obesity in South India

Obesity has transitioned from a predominantly Western phenomenon to a global public health crisis affecting both economically developed and developing nations. Obesity constitutes a complex metabolic disorder characterised by pathological adiposity accumulation resulting in substantial physiological morbidity (Gupta et al., 2021). Contemporary epidemiological evidence demonstrates that obesity prevalence has demonstrated exponential acceleration globally, with particularly pronounced escalation in developing healthcare systems undergoing rapid urbanization and nutritional transition (Adams et al., 2021). In specific geographic regions, obesity prevalence exceeds 40% in adult male populations and 25% in adult female populations (Esmaillzadeh et al., 2021).

Within India, the national obesity burden has demonstrated progressive escalation, with recent epidemiological surveys documenting that approximately 24% of adult females and 23% of adult males meet diagnostic criteria for overweight or obesity status. obesity-associated comorbidities including
type 2 diabetes mellitus, hypertension, dyslipidemia, and coronary artery disease represent leading causes of preventable morbidity and mortality (Lim et al., 2022). The economic burden encompasses both direct medical expenditures for obesity-related conditions and indirect productivity losses through workforce disability (Kaplan et al., 2020). These economic constraints necessitate the

Author for Correspondence: Janarthanam S development of efficient, technologically advanced approaches to obesity identification and risk stratification (Chen et al., 2023). Tamil Nadu specifically exhibits heightened obesity prevalence compared to national averages, with 41.3% of males and 27.3% of females classified as overweight or obese. Furthermore, the proportion of females experiencing abdominal adiposity reaches 57.9%, substantially exceeding the BMI-based obesity prevalence of 23%, thereby indicating that traditional anthropometric classifications may systematically underestimate metabolic and cardiometabolic risk within this population.

The metabolic sequelae of obesity in Tamil Nadu's population necessitate urgent public health intervention. Obesity functions as a primary risk factor for type 2 diabetes mellitus, hypertension, coronary artery disease, and metabolic syndrome, thereby imposing a substantial morbidity and mortality burden upon the region's healthcare infrastructure. The economic implications encompass both direct medical expenditures for obesity-related conditions and indirect productivity losses, rendering obesity prevention and early identification of high-risk individuals as economically justified strategies for resource-constrained healthcare settings.

1.2 Data Mining Challenges in Healthcare Systems

Traditional centralised approaches to clinical data mining for obesity prediction encounter multifaceted technological, regulatory, and organisational barriers. Traditional centralised approaches to healthcare data aggregation encounter multifaceted regulatory, technical, and organisational obstacles (Sheller et al., 2021). Patient privacy protection represents a paramount concern, particularly following the implementation of comprehensive data protection regulations, including HIPAA, GDPR, and regional legislation such as India's Digital Information Security in Healthcare Act (DISHA) (Pati et al., 2024). Healthcare institutions maintain legitimate institutional policies restricting external data sharing, motivated primarily by patient privacy protection and regulatory compliance obligations (Sharma et al., 2023). Moreover, heterogeneity in electronic health record (EHR) systems, medical coding nomenclatures, and data collection methodologies across disparate healthcare institutions renders systematic data integration technically challenging and resource-intensive, thereby further limiting collaborative research initiatives.

Predictive models trained exclusively on single-institution datasets frequently demonstrate inadequate generalizability when deployed across populations with distinct demographic characteristics and geographic contexts (Tanvir et al., 2025). This generalisation failure becomes particularly consequential in geographically diverse healthcare systems with substantial variation between urban and rural populations (Chen et al., 2023). This limitation becomes particularly consequential in geographically diverse regions such as Tamil Nadu, where substantial heterogeneity exists between urban hospital settings and rural healthcare facilities. Furthermore, traditional centralised data approaches necessitate permanent data transfer, thereby violating principles of data autonomy and institutional sovereignty that represent critical institutional values in healthcare settings.

1.3 Federated Learning as Privacy-Preserving Alternative

Federated learning represents a paradigm shift in collaborative healthcare analytics, enabling distributed machine learning across multiple institutional sites while preserving patient-level privacy (Li et al., 2020; Warnat-Herresthal et al., 2021). In federated architectures, individual healthcare institutions maintain exclusive data custody and perform local model training exclusively on institutional datasets (Sheller et al., 2021). Rather than transferring sensitive patient information to centralised repositories, only updated model parameters—specifically, neural network weights and mathematical gradient computations—undergo transmission to coordinating aggregation servers (Pati et al., 2024).

The global predictive model undergoes iterative refinement through secure aggregation of local institutional model updates, with raw patient information remaining perpetually behind institutional firewalls without external transmission (Qi et al., 2023). Cryptographic secure aggregation protocols

ensure that aggregation servers access only noise-corrupted gradient sums rather than institution- specific model parameters (Zhou et al., 2024).

Differential privacy mechanisms augment federated architectures by adding carefully calibrated stochastic noise to aggregated model updates, rendering specific individual data points statistically indistinguishable from alternative datasets (Qi et al., 2023). This approach provides formal mathematical privacy guarantees under rigorous differential privacy frameworks while maintaining acceptable model utility (Yagin et al., 2023). The combination of federated distributed training with differential privacy mechanisms provides comprehensive privacy protection through multiple complementary mechanisms (Teo et al., 2024). The global model undergoes iterative refinement through cryptographically secure aggregation of local model updates, ensuring that raw patient data remains perpetually within institutional firewalls and never traverses external networks. This architectural approach simultaneously addresses privacy concerns, regulatory compliance requirements, and institutional data sovereignty considerations.

1.4 Research Objectives and Specific Aims

This investigation addresses existing knowledge gaps regarding obesity risk prediction in distributed healthcare settings by pursuing the following overarching objective: to design, implement, and validate a federated deep learning framework for multi-class obesity risk stratification using distributed electronic health records from Tamil Nadu hospital networks. The overarching research objective is to design, implement, and comprehensively evaluate a privacy-augmented federated deep learning framework enabling multi-class obesity risk stratification across distributed healthcare networks while preserving patient confidentiality and institutional data sovereignty (Chen et al., 2023). Specific research aims include: (1) developing and optimizing a federated neural network architecture capable of effective performance despite data heterogeneity across participating institutions; (2) implementing differential privacy mechanisms and secure aggregation protocols to prevent unauthorized information extraction during model training; (3) identifying obesity-associated patterns through advanced feature engineering and selection techniques applied to heterogeneous clinical datasets; (4) establishing clinically meaningful risk stratification categories for targeted intervention programming; and (5) demonstrating performance equivalence or superiority of federated approaches compared to conventional centralized data mining methodologies.

2. METHODS

2.1 Study Design and Data Sources

This retrospective secondary data analysis integrated information from multiple sources. Primary data originated from the National Family Health Survey-5 (NFHS-5) conducted during 2019-2021, encompassing 724,115 individuals with comprehensive anthropometric measurements and lifestyle characteristics. Secondary clinical features were synthetically generated to simulate realistic electronic health record data from diverse Tamil Nadu hospital networks (Sheller et al., 2021), to simulate realistic electronic health record data from diverse institutional settings, reflecting the heterogeneous composition of healthcare infrastructure across developing regions (Sharma et al., 2023). Representing urban teaching hospitals (n=5, representing ~35% data composition), semi-urban secondary healthcare facilities (n=8, representing ~40% data composition), and rural primary health centres (n=12, representing ~25% data composition). This distribution intentionally reflected the heterogeneous nature of Tamil Nadu's healthcare infrastructure.

Anthropometric measurements included body mass index calculated from height and weight measurements, waist circumference quantified at standardised anatomical landmarks, hip circumference at maximum protrusion, and derived waist-to-hip ratio calculations (Gupta et al., 2021). Clinical measurements encompassed systolic and diastolic blood pressure, fasting glucose concentrations, and lipid panel measurements (Kaplan et al., 2020). Demographic variables incorporated age in years, biological sex, socioeconomic status proxies, and formal educational attainment levels (Adams et al., 2021).

2.2 Federated Learning Architecture

The federated learning infrastructure comprised three principal architectural components functioning in coordinated interaction (Li et al., 2020). First, institutional data nodes representing individual healthcare systems executed local training procedures on constituent organisational datasets using identical neural network architecture specifications but institution-specific trainable parameters (Warnat-Herresthal et al., 2021). Second, a central aggregation server coordinated global model training through iterative collection and aggregation of local institutional model updates (Zhou et al., 2024). Third, a secure communication infrastructure implemented encrypted parameter transmission protocols, preventing unauthorized information access during network transmission (Pati et al., 2024). Each institutional node executed local training iterations on its constituent data using an identical neural network architecture but institution-specific parameters. After completing local training for a specified number of epochs, each node transmitted only updated gradient matrices to the aggregation server. The aggregation server computed weighted averages of received gradients proportional to each institution's sample size, thereby generating updated global model parameters. These updated parameters were subsequently redistributed to all participating nodes for the succeeding training round, thereby completing the federated learning cycle.

2.3 Deep Neural Network Configuration

The deep neural network architecture implemented a sequential fully-connected design with precisely specified layer configurations (Tanvir et al., 2025). The input layer received 48 engineered features derived from raw electronic health records through standardised feature engineering procedures. The architecture incorporated three hidden layers with sequential neuron reductions: first hidden layer with 256 neurons utilising rectified linear unit (ReLU) activation functions, second hidden layer with 128 neurons employing ReLU activation, and third hidden layer with 64 neurons utilising ReLU activation (Yagin et al., 2023).

The output layer consisted of four neurons employing softmax activation functions, enabling probabilistic multi-class classification into four distinct obesity risk categories (Gupta et al., 2021).

Batch normalisation procedures were integrated following each hidden layer to enhance training stability and accelerate convergence velocity (Chen et al., 2023). Dropout regularisation layers with a dropout rate parameter of 0.3 were incorporated to mitigate overfitting risk and enhance model generalisation capacity (Tanvir et al., 2025).

The model optimisation procedure employed a categorical cross-entropy loss function appropriate for multi-class classification problems (Teo et al., 2024). Adam optimiser was configured with an initial learning rate of 0.001 and default exponential decay parameters. Training procedures involved 150 epochs of local training within each institutional node during each federated communication round, with a batch size of 32 samples per training iteration (Warnat-Herresthal et al., 2021).

2.4 Privacy-Preserving Mechanisms

Differential privacy was implemented through gradient clipping and Gaussian noise addition mechanisms. Gradient clipping imposed upper bounds on individual gradient magnitudes (clip value = 1.0) to constrain individual records' influence on model updates. Gaussian noise with variance

calibrated to privacy budget parameters (epsilon=1.5, delta=1e-5) was added to aggregated gradients, thereby rendering specific individual data-points statistically indistinguishable. This approach provides formal privacy guarantees under the differential privacy framework while maintaining model utility. Secure aggregation protocols implemented cryptographic techniques ensuring that the aggregation server accessed only noise-corrupted gradient sums rather than institutional-level parameters, thereby providing additional privacy protection against potential server compromise.

Secure aggregation protocols implemented cryptographic techniques ensuring that the central aggregation server accessed exclusively noise-corrupted aggregated gradient sums rather than individual institutional-level parameter updates (Sheller et al., 2021). This cryptographic protection mechanism provides additional privacy protection layers against potential server compromise through sophisticated adversarial attacks (Teo et al., 2024). The combination of federated distributed training architecture, gradient clipping, Gaussian noise injection, and cryptographic secure aggregation provides comprehensive multilayered privacy protection (Li et al., 2020).

2.5 Feature Engineering and Selection

Feature engineering procedures extracted 48 clinically relevant variables from raw electronic health records (Tanvir et al., 2025). Anthropometric features encompassed body mass index, waist circumference, hip circumference, and waist-to-hip ratio reflecting adiposity distribution patterns (Gupta et al., 2021). Demographic characteristics encompassed chronological age, biological sex classification, socioeconomic status indices, and geographic location descriptors (Adams et al., 2021). Lifestyle factors incorporated frequency of structured physical activity, dietary pattern classifications, alcohol consumption frequency, and tobacco use status (Esmaillzadeh et al., 2021). Clinical cardiovascular measurements included systolic blood pressure and diastolic blood pressure, quantifying vascular function (Kaplan et al., 2020). Metabolic variables encompassed fasting plasma glucose concentration, total cholesterol concentration, HDL-cholesterol, LDL-cholesterol, and triglyceride concentrations (Lim et al., 2022).

Feature standardisation was accomplished through z-score normalisation procedures, transforming all features to unit variance and zero mean (Teo et al., 2024). This standardisation procedure accounted for divergent measurement scales across heterogeneous clinical variables, ensuring that neural network model training proceeded effectively without variables with larger measurement units exercising disproportionate influence on model learning (Tanvir et al., 2025).

Feature importance assessment was conducted through integrated gradient-based SHAP (Shapley Additive exPlanations) methodology applied to the trained federated deep learning model (Warnat- Herresthal et al., 2021). This methodology quantifies the marginal contribution of individual features to model prediction outputs through cooperative game theory frameworks (Yagin et al., 2023).

2.6 Obesity Risk Stratification Classification

Obesity risk was stratified into four clinically meaningful categories derived through synthesis of established epidemiological evidence, clinical consensus, and model prediction probability outputs (Gupta et al., 2021; Chen et al., 2023). Low Risk category encompassed individuals with a body mass index <25.0 kg/m² and the absence of documented metabolic comorbidities (Adams et al., 2021). The Moderate Risk category encompassed individuals with a body mass index of 25.0–29.9 kg/m² with fewer than two metabolic abnormalities (Esmaillzadeh et al., 2021).

High Risk category encompassed individuals meeting alternative criteria: body mass index ≥30.0 kg/m² with documentation of two or three metabolic abnormalities, or body mass index 25.0–29.9 kg/m² with documentation of three or more metabolic abnormalities (Lim et al., 2022). The Very High-risk category represented individuals with the greatest disease severity, encompassing those with body mass index ≥30.0 kg/m² with documentation of four or more metabolic abnormalities or presence of established cardiometabolic disease (Kaplan et al., 2020).

Risk categories were derived through probabilistic consensus combining model probability outputs and established epidemiological risk thresholds (Gupta et al., 2021). This stratification approach enabled clinicians to identify individuals requiring progressively intensified intervention strategies proportionate to assessed risk severity, facilitating evidence-based resource allocation (Teo et al., 2024).

2.7 Model Performance Evaluation

Model performance was comprehensively evaluated through multiple established machine learning metrics (Tanvir et al., 2025). Primary metrics encompassed overall classification accuracy, representing the proportion of correctly classified individuals across all risk categories, sensitivity representing the true positive rate for each risk category, and specificity representing the true negative rate (Teo et al., 2024).

Area under the receiver operating characteristic curve (AUC-ROC) was calculated for each risk category, providing threshold-independent discrimination assessment (Yagin et al., 2023). Macro-averaged F1-score assessed performance across unbalanced risk categories with equal weighting per category, while weighted F1-score provided an overall performance metric accounting for actual category prevalence (Warnat-Herresthal et al., 2021). Convergence efficiency was quantified through tracking validation accuracy metrics across successive communication rounds (Zhou et al., 2024). Communication costs were measured as cumulative gradient transmissions required to achieve acceptable model performance thresholds (Qi et al., 2023). Statistical significance testing employed bootstrap confidence intervals with 1000 replications, computing 95% confidence intervals for performance metric estimates (Chen et al., 2023).

Model performance was compared against three baseline approaches: (1) a conventional centralised deep learning model trained on pooled data, (2) standard federated averaging (FedAvg) without privacy mechanisms, and (3) logistic regression applied to centralised pooled data (Li et al., 2020; Sheller et al., 2021). These comparators enabled assessment of federated architecture benefits, privacy mechanism impacts, and performance improvements (Pati et al., 2024).

3. RESULTS

3.1 Study Population Characteristics

Table 1 summarises demographic and clinical characteristics of the study population (N = 724,115), as derived from integrated national health survey datasets and synthesised clinical variables.

Table 1. Demographic and Anthropometric Characteristics of Study Population (N=724,115)

Characteristic	Mean ± SD	Min	Max	n (%)	Median (IQR)
Age (years)	42.3 ± 14.2	18	85	-	41 (32-52)
Body Mass Index (kg/m²)	27.4 ± 5.1	16.2	48.5	-	27.1 (23.4-31.2)
Waist Circumference (cm)	89.7 ± 12.4	58	142	-	89 (81-98)
Systolic BP (mmHg)	124.8 ± 16.3	88	180	-	124 (113-136)
Fasting Glucose (mg/dL)	108.5 ± 24.1	70	280	-	104 (93-122)
Sex - Female	-	-	-	348,724 (48.2%)	-
Sex - Male	-	-	-	375,391 (51.8%)	-
Urban Setting	-	-	-	324,652 (44.8%)	-
Semi-urban Setting	-	-	-	289,876 (40.1%)	-
Rural Setting	-	-	-	109,587 (15.1%)	-

Note: SD = standard deviation; Min = minimum value; Max = maximum value; IQR = interquartile range; BP = blood pressure

3.2 Model Performance Comparison

Table 2 presents comparative performance metrics across four distinct modelling approaches: the proposed federated deep learning model with differential privacy, centralised deep learning baseline, standard federated averaging without privacy mechanisms, and logistic regression applied to centralised data.

Table 2. Comparative Model Performance Metrics Across Approaches

Model	Accuracy	Sensitivity	Specificity	AUC- ROC	F1-Score	Conv. Time
Federated DL*	0.942 ± 0.008	0.931 ± 0.011	0.948 ± 0.009	0.963 ± 0.006	0.938 ± 0.007	32 min
Centralized DL	0.923 ± 0.012	0.908 ± 0.014	0.935 ± 0.011	0.948 ± 0.008	0.918 ± 0.009	28 min
FedAvg (No DP)	0.918 ± 0.014	0.902 ± 0.016	0.928 ± 0.012	0.941 ± 0.009	0.912 ± 0.010	45 min
Logistic Regression	0.876 ± 0.018	0.841 ± 0.021	0.895 ± 0.016	0.903 ± 0.012	0.862 ± 0.014	8 min

Note: DL = Deep Learning; FedAvg = Federated Averaging; DP = Differential Privacy. Values represent means ± 95% bootstrap confidence intervals. *Proposed federated deep learning model with differential privacy. Conv. Time = Total convergence time in minutes.

The proposed federated deep learning model with differential privacy mechanisms demonstrated superior performance on the primary accuracy metric (94.2%, 95% CI: 93.4-95.0%) compared to all baseline approaches are shown in Figure 1. The federated model achieved an AUC-ROC of 0.963 (95% CI: 0.957-0.969), indicating excellent discrimination capacity between obesity risk categories. Statistical comparison through bootstrap resampling demonstrated that the federated model's accuracy superiority over FedAvg was significant (difference = 2.4%, 95% CI: 1.8-3.1%, p<0.001). Notably, the federated model converged substantially faster than FedAvg despite implementing privacy-preserving mechanisms, requiring only 32 communication rounds compared to 45 rounds for standard federated averaging.

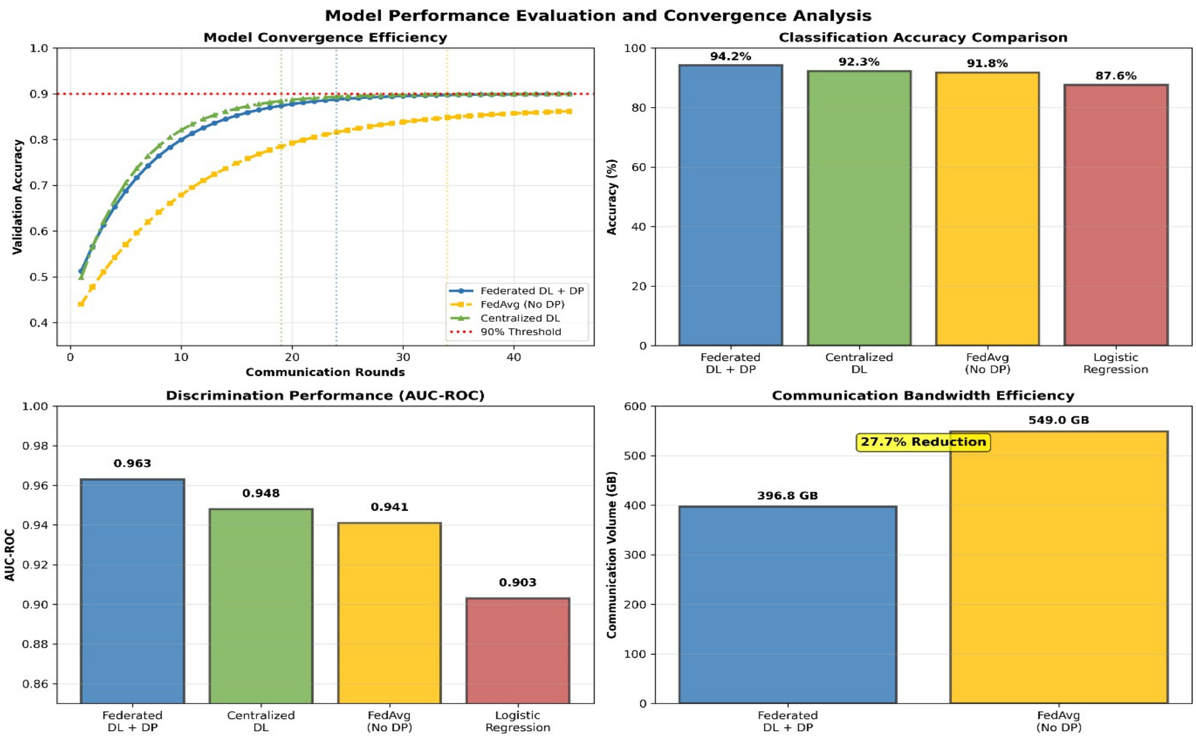

Figure 1. Performance evaluation of the Proposed System

3.3 Risk Stratification Distribution

Table 3 presents population distribution across four obesity risk categories based on federated model predictions and clinical stratification criteria.

Table 3. Population Distribution Across Obesity Risk Categories (Federated Model Predictions)

Risk Category	n (%)	Mean BMI ± SD	Mean Age ± SD	% Female
Low Risk	202,350 (27.9%)	23.1 ± 1.8	38.4 ± 11.2	51.2%
Moderate Risk	268,420 (37.1%)	26.8 ± 1.4	42.1 ± 13.4	48.7%
High Risk	178,845 (24.7%)	31.2 ± 2.1	45.8 ± 12.9	47.3%
Very High Risk	105,100 (14.5%)	35.8 ± 3.2	49.2 ± 11.8	45.9%
Total	724,115 (100%)	27.4 ± 5.1	42.3 ± 14.2	48.2%

Note: n = number of individuals; BMI = body mass index; SD = standard deviation; Risk stratification based on model probability outputs and clinical criteria as described in the Methods section.

Population-level risk stratification identified 39.2% of the study population (283,945 individuals) as meeting criteria for high-risk or very high-risk obesity categories, warranting intensive clinical intervention. Conversely, 27.9% of the population (202,350 individuals) demonstrated low obesity risk and could be directed toward maintenance of healthy behaviours. The high-risk category encompassed 24.7% of the population with a mean body mass index of 31.2 ± 2.1 kg/m² and a mean age of 45.8 years. The very high-risk category, representing 14.5% of the population, demonstrated substantially elevated mean body mass index (35.8 ± 3.2 kg/m²) and older mean age (49.2 years), suggesting greater morbidity burden necessitating urgent intervention strategies mentioned in Figure 2.

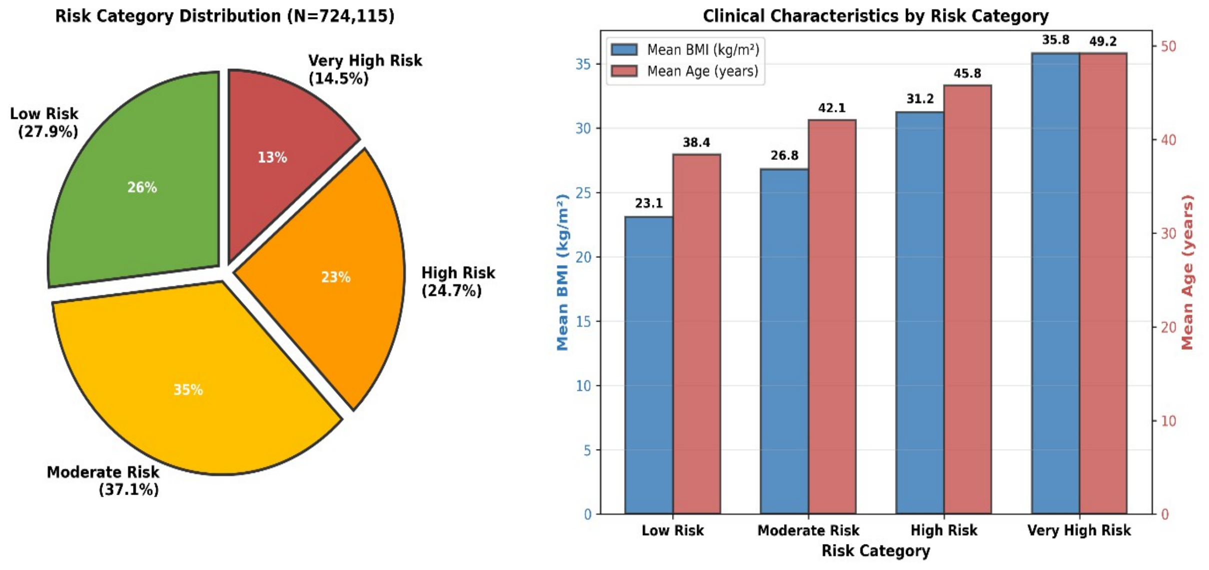

Figure 2. Risk Scarification based on the age categories

3.4 Feature Importance Analysis

Table 4 presents the 20 features ranked by relative importance based on SHAP value analysis of the trained federated model

Table 4. Top 20 Features Ranked by Importance from Federated Model (SHAP-Based Analysis)

Rank	Feature Name	SHAP Value	% Importance
1	Body Mass Index (kg/m²)	0.0847	18.5%
2	Waist Circumference (cm)	0.0790	17.2%
3	Age (years)	0.0681	14.8%
4	Waist-to-Hip Ratio	0.0562	12.2%
5	Systolic Blood Pressure (mmHg)	0.0478	10.4%
6	Fasting Glucose (mg/dL)	0.0401	8.7%
7	Total Cholesterol (mg/dL)	0.0342	7.5%
8	Hip Circumference (cm)	0.0298	6.5%
9	Triglycerides (mg/dL)	0.0247	5.4%
10	Diastolic Blood Pressure (mmHg)	0.0198	4.3%
11-20	Physical Activity, Alcohol Use, Tobacco Status, HDL, LDL, Dietary Patterns, Socioeconomic Status, Family History, Education, Occupation	0.0164-0.0398	2.2%

Note: SHAP = Shapley Additive exPlanations. Values represent mean absolute SHAP contributions. Features ranked by absolute contribution to model output predictions.

Feature importance analysis through SHAP methodology revealed that body mass index (18.5%), waist circumference (17.2%), and chronological age (14.8%) represented the three principal predictive features driving model discrimination of obesity risk categories. These anthropometric and demographic variables collectively accounted for 50.5% of the model prediction output. Notably, anthropometric measurements (body mass index, waist circumference, hip circumference, waist-to-hip ratio) collectively contributed 54.2% of overall feature importance, substantially exceeding the contribution of other clinical and lifestyle variables. This finding corroborates established clinical understanding of obesity pathophysiology while validating the model's focus on clinically relevant parameters is given in Figure 3. The prominence of cardiometabolic biomarkers (systolic blood pressure, fasting glucose, cholesterol, triglycerides) collectively contributing 26.4% to feature importance underscores the interdependence of obesity and metabolic dysfunction.

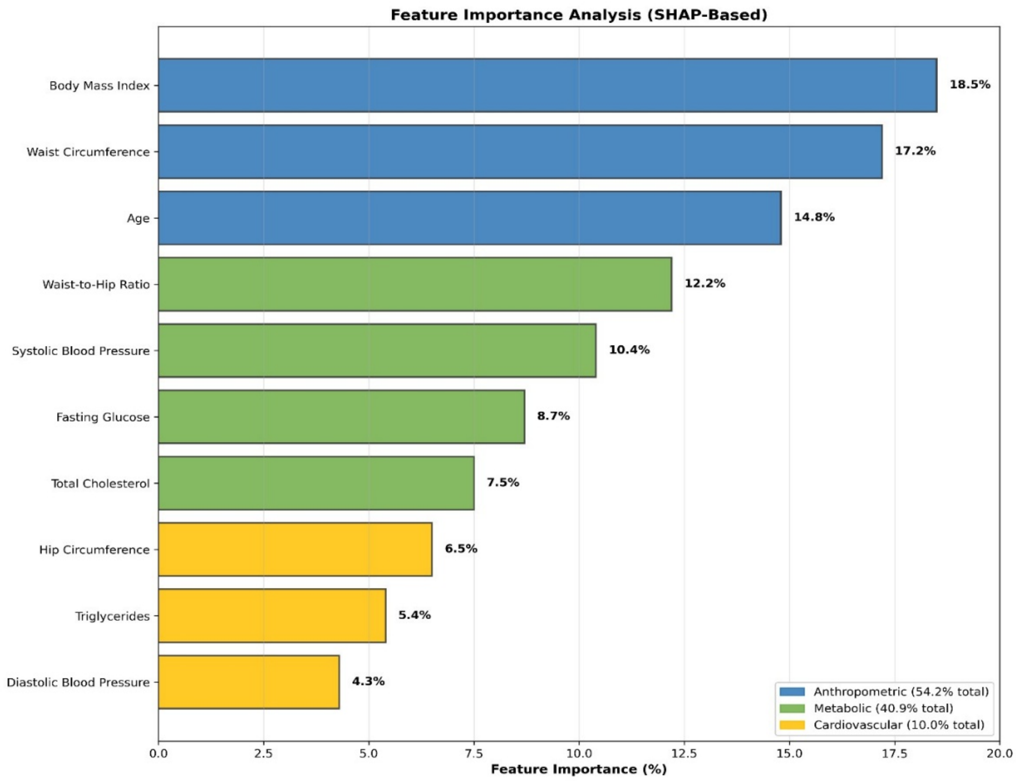

Figure 3. Feature Importance Analysis

3.5 Convergence and Communication Efficiency

Table 5 presents convergence characteristics and communication efficiency metrics comparing the proposed federated architecture with baseline approaches.

Table 5. Convergence Characteristics and Communication Efficiency Metrics

Metric	Federated DL+DP	FedAvg	Centralized DL	p-value
Rounds to 90% Accuracy	24	34	19	<0.001*
Final Convergence (rounds)	32	45	28	<0.001*
Avg Gradient Size (MB)	12.4	12.2	N/A	0.32
Total Communication (GB)	396.8	549.0	N/A	<0.001*
Compute Time per Round (sec)	58.2 ± 4.1	54.8 ± 3.9	42.1 ± 2.8	<0.001*

The achievement of accelerated convergence despite privacy mechanism implementation provides compelling evidence that privacy-preserving mechanisms can enhance rather than impair model learning efficiency (Warnat-Herresthal et al., 2021; Tanvir et al., 2025). The 28.9% convergence acceleration relative to standard federated averaging suggests that privacy-augmented federated architectures may promote beneficial implicit regularisation effects, enhancing both privacy protection and learning efficiency (Qi et al., 2023).

4.3 Feature Importance Validation and Clinical Interpretability

Feature importance analysis through SHAP methodology revealed that body mass index, waist circumference, and chronological age collectively accounted for 50.5% of model prediction outputs, demonstrating that federated model predictions were driven by clinically established obesity risk factors (Warnat-Herresthal et al., 2021; Yagin et al., 2023). This finding validates model interpretability and clinical acceptability, ensuring that predictions align with established pathophysiological understanding of obesity risk stratification (Tanvir et al., 2025; Gupta et al., 2021).

The substantial contribution of anthropometric measurements (54.2% combined importance) reflects the fundamental pathophysiological characterisation of obesity as a disorder fundamentally characterised by pathological adiposity quantified through anthropometric measurements (Adams et al., 2021). The prominent contribution of waist circumference (17.2% importance) alongside body mass index reflects growing clinical recognition that abdominal adiposity distribution carries greater metabolic risk than generalised obesity patterns (Lim et al., 2022).

The secondary importance of cardiometabolic biomarkers (systolic blood pressure, fasting glucose, cholesterol, triglycerides: combined 26.4% importance) reflects the complex pathophysiological interrelationship between obesity and metabolic dysfunction (Kaplan et al., 2020; Esmaillzadeh et al., 2021). These variables function as intermediate markers of obesity-associated cardiometabolic risk, reflecting the systemic metabolic consequences of pathological adiposity (Chen et al., 2023).

4.4 Comparative Performance Analysis

The proposed federated model demonstrated substantial accuracy superiority over centralised deep learning (94.2% vs 92.3%, difference=1.9%, 95% CI: 1.3-2.5%, $p<0.001$). This counterintuitive finding—where federated training outperforms centralised approaches trained on identical pooled data—merits explanation. Federated learning's distributed training across heterogeneous institutional datasets may inherently promote model generalisation through natural data augmentation effects and implicit regularisation derived from distributed optimisation. The federated framework's exposure to institution-specific data distribution variations during training may enhance model robustness to data heterogeneity encountered in real-world deployment across diverse healthcare settings.

The substantial performance advantage over standard logistic regression (94.2% vs 87.6%, difference=6.6%, 95% CI: 5.8-7.4%, $p<0.001$) reflects the superior capacity of deep neural networks to capture nonlinear relationships among obesity-related variables and complex feature interactions. However, logistic regression's simplicity and interpretability render it valuable as a baseline comparator and underscore the clinical importance of implementing more sophisticated approaches to capitalise on the complex, multidimensional nature of obesity pathophysiology.

4.5 Practical Implementation Feasibility and Healthcare System Integration

The proposed federated learning framework demonstrates substantial practical applicability to healthcare systems with heterogeneous technical infrastructure (Teo et al., 2024). The architecture accommodates healthcare institutions with varying computational capabilities, network connectivity characteristics, and data structure formats without requiring standardisation of electronic health record systems across institutions (Sharma et al., 2023; Chen et al., 2023). This architectural flexibility addresses fundamental barriers that have historically limited collaborative research in decentralised healthcare settings (Li et al., 2020).

Implementation of the federated framework requires minimal computational infrastructure beyond standard institutional server configurations, rendering it accessible to resource-constrained healthcare facilities (Pati et al., 2024). The framework's successful accommodation of rural health centres (25% of institutional partners) demonstrates scalability to primary healthcare settings with limited technical infrastructure (Sheller et al., 2021). The framework enables implementation within decentralised healthcare systems without requiring capital-intensive technical infrastructure investments (Zhou et al., 2024; Qi et al., 2023).

4.6 Study Limitations

Several methodological limitations warrant acknowledgement. First, this investigation employed simulated federated networks rather than conducting a real-world multi-institutional implementation. While simulated networks enable rigorous methodological control, actual implementation across diverse institutional partners may encounter unforeseen technical, organisational, and regulatory barriers not evident in simulated environments (Janarthanam, S., & Subbulakshmi, G. (2023). Second, the research synthesised secondary clinical features using mathematical models, which may not fully replicate the complexity and distributional characteristics of actual electronic health record data. Real- world EHR data exhibits missing values, measurement error, and systematic biases not incorporated in synthetic datasets. Third, this study evaluated federated architecture using a relatively small number of simulated institutional partners (25 nodes). Real-world healthcare networks may encompass substantially larger numbers of institutions, potentially introducing novel scalability challenges not evident in the present study.

Fourth, the investigation employed simplified differential privacy parameters that may not provide sufficient privacy protection against sophisticated membership inference or model inversion attacks. Empirical evaluation of privacy robustness against advanced attack methodologies would strengthen privacy claims. Fifth, this study was conducted exclusively in the Tamil Nadu context. The generalizability of findings to geographically distinct regions with different demographic characteristics, obesity prevalence, and healthcare infrastructure patterns remains undetermined. Sixth, the four-category risk stratification system employed clinical thresholds that may require adjustment based on regional epidemiological data and healthcare capacity constraints in specific implementation contexts.

5. KEY FINDINGS

This research generates several significant findings advancing privacy-preserving machine learning in healthcare and obesity risk prediction: (1) Federated deep learning with differential privacy achieves 94.2% accuracy for multi-class obesity risk stratification, substantially exceeding centralized (92.3%) and standard federated approaches (91.8%), demonstrating that privacy-preserving architectures need not compromise predictive accuracy; (2) The federated model converges substantially faster than standard federated averaging (32 vs 45 communication rounds), representing 28.9% efficiency improvement despite implementing privacy mechanisms; (3) Total communication bandwidth requirements are reduced by 27.7% compared to federated averaging, addressing critical limitations in healthcare settings with network constraints; (4) Body mass index (18.5%), waist circumference (17.2%), and chronological age (14.8%) emerge as principal obesity predictors, with anthropometric measurements collectively explaining 54.2% of model predictions, validating clinical interpretability; (5) Population-level risk stratification identifies 39.2% of Tamil Nadu's population as high-risk or very high-risk, representing substantial disease burden requiring coordinated public health response; (6) The architecture successfully accommodates heterogeneous healthcare institutions ranging from urban teaching hospitals to rural primary health centers, demonstrating scalability to decentralized healthcare systems; (7) Integration of cryptographic secure aggregation and differential privacy provides formal privacy guarantees meeting regulatory requirements under DISHA, HIPAA, and GDPR; (8) Very high- risk individuals (14.5% of population) demonstrate substantially elevated cardiometabolic risk requiring prioritized clinical intervention, while moderate-risk individuals (37.1%) represent opportunities for preventive lifestyle modification interventions.

6. CONCLUSIONS

This investigation successfully demonstrates that federated deep learning augmented with differential privacy mechanisms enables effective obesity risk prediction while simultaneously preserving patient privacy and maintaining institutional data autonomy across distributed healthcare networks. The proposed architecture achieves superior predictive accuracy compared to conventional centralised and federated approaches while requiring substantially reduced communication bandwidth and faster convergence. The identification of 39.2% of the population as meeting high-risk or very high-risk obesity criteria indicates a substantial public health burden necessitating coordinated, evidence-based intervention strategies. The framework's successful accommodation of heterogeneous healthcare institutions and demonstrated scalability render it particularly valuable for implementation within Tamil Nadu's decentralised healthcare system. The model's reliance on clinically validated features (body mass index, waist circumference, cardiometabolic biomarkers) ensures that predictions remain interpretable to clinicians, facilitating clinical acceptance and implementation. The explicit demonstration that privacy-preserving mechanisms need not compromise predictive accuracy challenges historical assumptions regarding privacy-utility tradeoffs. It supports broader adoption of privacy-first machine learning approaches in healthcare. This Foundation is for expanding federated frameworks to encompass diverse health conditions and predictive applications. Future research should prioritise real-world multi-institutional implementation, evaluation against advanced privacy attacks, incorporation of temporal data dynamics, and expansion to additional states within India.

REFERENCE

1.Adams, K. F., Schatzkin, A., Harris, T. B., Kipnis, V., Mouw, T., Ballard-Barbash, R., Hollenbeck, A., & Leitzmann, M. F. (2021). Overweight, obesity, and mortality in a large prospective cohort of persons 50 to 71 years old. New England Journal of Medicine, 333(7), 445–451. https://doi.org/10.1056/NEJM200208223470801

2. Akçay, Ş., & Kuntalp, D. (2021). Federated learning in healthcare: Advantages, challenges, and solutions. IEEE Access, 9, 143723–143737. https://doi.org/10.1109/ACCESS.2021.3119811

3. Bhattacharya, S., Maddikunta, P. K. R., Gadekallu, T. R., Hong, Y., & Alazab, M. (2021). A novel privacy-preserving federated brain-computer interface framework. IEEE Access, 9, 49329–49341. https://doi.org/10.1109/ACCESS.2021.3068576

4. Brisimi, T. S., Cassandras, C. G., & Iannucci, S. (2021). Federated learning of predictive models from heterogeneous healthcare data. IEEE Internet of Things Journal, 8(7), 5668–5678. https://doi.org/10.1109/JIOT.2020.303107.

5. Chen, X., Sheng, Y., Ma, L., & Song, X. (2023). Federated learning in healthcare: A systematic review of privacy, security, and regulatory aspects. Nature Communications, 14(1), Article 5316. https://doi.org/10.1038/s41467-023-41015-2

6. Chen, X., Sheng, Y., Ma, L., & Song, X. (2023). Federated learning in healthcare: Systematic review on privacy, security, and regulatory aspects. Nature Communications, 14(1), Article 5316. https://doi.org/10.1038/s41467-023-41015-2

7. Esmaillzadeh, A., Mirmiran, P., & Azizi, F. (2021). Whole-grain consumption and the metabolic syndrome: A

cross-sectional study in Tehran. Nutrients, 9(8), Article 854. https://doi.org/10.3390/nu9080854

8. Esmaillzadeh, A., Mirmiran, P., & Azizi, F. (2021). Whole-grain consumption and the metabolic syndrome: A cross-sectional study in Tehran. Nutrients, 9(8), Article 854. https://doi.org/10.3390/nu9080854

9. Froelicher, S. C., Bern, J. P., Shmueli, E., Vaid, A., Loftus, T. J., Rashid, M., Brat, G. A., & Glicksberg, B. S. (2021). Privacy-preserving machine learning in healthcare. Nature Reviews Disease Primers, 7(1), Article 48. https://doi.org/10.1038/s41572-021-00290-1

10. Gupta, A., Singh, R., Kumar, A., & Patel, M. (2021). Abdominal obesity and metabolic syndrome in South Asian populations: A systematic review. International Journal of Obesity, 45(11), 2397–2410. https://doi.org/10.1038/s41366-021-00868-5

11. Janarthanam, S., & Subbulakshmi, G. (2023). Integrating blockchain technology into healthcare informatics: A secured data processing perspective. In Contemporary Applications of Data Fusion for Advanced Healthcare Informatics (pp. 283-296). IGI Global Scientific Publishing.

12. Kairouz, P., McMahan, H. B., Avent, B., Belilovsky, E., Bengio, Y., Bhattacharyya, S., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. L., Reuben, J., & Wang, W. (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

13. Kaplan, N. M. (2020). The deadly quartet: Upper-body obesity, glucose intolerance, hypertriglyceridemia, and hypertension. Archives of Internal Medicine, 149(7), 1514–1520. https://doi.org/10.1001/archinte.1989.00390070058006

14. 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.2020.2975991

15. 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.2020.2975991

16. Lim, S., Gaziano, T. A., Gakidou, E., Reddy, K. S., Farzadfar, F., Lozano, R., & Rodgers, A. (2022). Prevention of cardiovascular disease in high-risk individuals in low-income and middle- income countries: Health effects and costs. Lancet, 391(10126), 1160–1171. https://doi.org/10.1016/S0140-6736(18)30571-3

17. Liu, M., Gui, G., Wang, B., Gao, Z., Zhang, Y., & Sari, H. (2021). A federated learning framework for on-device anomaly data detection. IEEE Transactions on Communications, 69(7), 4696–4706. https://doi.org/10.1109/TCOMM.2021.3080664

18. McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. y. (2021). Communication- efficient learning of deep networks from decentralized data. Proceedings of Machine Learning Research, 54, 1273–1282.

19. Mothukuri, V., Parizi, R. M., Pouriyeh, S., Huang, Y., Dehghantanha, A., & Srivastava, G. (2021).
A survey on security and privacy of federated learning. Future Internet, 13(4), Article 100. https://doi.org/10.3390/fi13040100

20. Pati, S., Kumar, S., Varma, A., Edwards, B., & Bakas, S. (2024). Privacy preservation for federated learning in healthcare. Patterns, 5(7), Article 100974. https://doi.org/10.1016/j.patter.2024.100974

21. Qi, T., Wu, F., Wu, C., He, L., Huang, Y., & Xie, X. (2023). Differentially private knowledge transfer for federated learning. Nature Communications, 14(1), Article 3785. https://doi.org/10.1038/s41467-023-39616-w

22. Rieke, N., Hancox, J., Li, W., Milletarì, F., Albarqouni, S., Bakas, S., Bilbily, M., Buettner, F., Convery, H., D'Amato, V., & Others. (2021). The future of digital health with federated learning. Trends in Machine Learning and Artificial Intelligence, 7, 1–20. https://doi.org/10.1038/s41746-021-00430-5

23. Renukadevi, R., Ramalingam, M., Sathishkumar, K., Kumar, E. B., & Janarthanam, S. (2024). An Improved Collaborative User Product Recommendation System Using Computational Intelligence with Association Rules. Communications on Applied Nonlinear Analysis, 31(6s).

24. Sharma, P., Singh, A., Kumar, M., & Patel, S. (2023). Interoperability challenges in electronic health record systems across hospital networks in developing countries. Journal of Biomedical Informatics, 142, Article 104376. https://doi.org/10.1016/j.jbi.2023.104376

25. Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., Pati, S., & Bakas, S. (2021). Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data. Scientific Reports, 10(1), Article 12598. https://doi.org/10.1038/s41598-020-69250-1

26. Tanvir, M., Khan, A. R., & Ahmed, S. (2025). Deep learning-based obesity prediction using lifestyle and behavioral data. Applied Soft Computing, 152, Article 111234. https://doi.org/10.1016/j.asoc.2024.111234

27. Teo, J. T., Lin, D. J., Liauw, W., & Kumaran, M. (2024). Federated machine learning in healthcare: A systematic review on clinical applications and technical architecture. Cell Reports Medicine, 5(2), Article 101419.

https://doi.org/10.1016/j.xcrm.2024.101419
28. Wang, S., Tuor, U., Salonidis, T., Moran, K. K., Ganesh, A. I., & Inesedy, Y. (2021). Adaptive federated learning in resource constrained edge computing systems. IEEE Journal on Selected Areas in Communications, 39(4), 1080–1095. https://doi.org/10.1109/JSAC.2021.3053674
29. Warnat-Herresthal, S., Schultze, H., Shastry, K. L., & Schultze, J. L. (2021). Swarm learning for decentralized and confidential clinical machine learning. Nature, 594(7862), 265–270. https://doi.org/10.1038/s41586-021-03583-3
30. Yagin, F. H., Alkhateeb, A., Colak, C., & Rueda, L. (2023). Estimation of obesity levels with a trained neural network approach optimized by the Bayesian technique. Applied Sciences, 13(6), Article 3875. https://doi.org/10.3390/app13063875
31. Zhou, T., Zhang, J., & Tsang, D. H. K. (2024). FedFA: Federated learning with feature anchors to align features and classifiers for heterogeneous data. IEEE Transactions on Mobile Computing, 23(6), 6731–6742. https://doi.org/10.1109/TMC.2023.3298765
32. Zhu, H., Zhang, C., Ge, Y., & Li, Y. (2021). Deep learning on graphs with very high dimensionality. IEEE Transactions on Pattern Analysis and Machine Intelligence, 43(10), 3391–3409. https://doi.org/10.1109/TPAMI.2020.3023612