

FedSS: Shareability-Aware Privacy-Preserving Federated Learning for Sports Recognition

Jiahui Lin

University of Chinese Academy of Sciences

Jie Wu

wujie@hrbipe.edu.cn

Kyungil University

Ke Sun

University of Chinese Academy of Sciences

Yuanyuan Huang

Chengdu University of Information Technology

Article

Keywords: federated learning, differential privacy, sports image recognition

Posted Date: October 21st, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11163668/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

FedSS: Shareability-Aware Privacy-Preserving Federated Learning for Sports Recognition

Jiahui Lin¹, Jie Wu^{2*}, Ke Sun¹, Yuanyuan Huang^{3*}

¹Department of Physical Education, University of Chinese Academy of Sciences, Beijing, China.

²College of Art and Physical Education, Kyungil University, 50, Gamasil-gil, Hayang-eup, Gyeongsan-si, Gyeongbuk-do, Republic of Korea.

³School of Artificial Intelligence, Chengdu University of Information Technology, Chengdu, 611025, China.

*Corresponding author(s). E-mail(s): wujie@hrbipe.edu.cn;
iyyhuang@hotmail.com;

Contributing authors: linjh@ucas.ac.cn; sunke@ucas.ac.cn;

Abstract

Federated learning provides a promising approach for collaborative sports image recognition without directly sharing raw data across organizations. However, model updates may still expose sensitive information, motivating the use of differentially private federated learning (DP-FL). A major challenge is that heterogeneous sports data contain both transferable sport semantics and client-specific visual context. Conventional DP-FL processes these components uniformly, causing less transferable information to consume part of the limited clipping range and increasing the distortion of useful sport knowledge. To address this problem, we propose **FedSS**, a shareability-aware privacy-preserving federated learning framework for sports image recognition. FedSS introduces a Sport Semantic Shareability Module (SSSM) that evaluates stage-wise model components according to sport discriminativeness, cross-client invariance, and privacy robustness. The resulting shareability scores emphasize transferable sport-semantic information during private federation while retaining client-specific context locally. First-order gradient stability is further incorporated to improve the robustness of shared updates under privacy perturbation. Experiments on sport22 and sport100 under different heterogeneous and privacy settings show that FedSS consistently improves recognition accuracy and the privacy-utility trade-off over representative DP-FL methods. Further analyses demonstrate stronger

cross-client semantic consistency and reduced clipping distortion, supporting the effectiveness of shareability-aware private federation for sports image recognition.

Keywords: federated learning, differential privacy, sports image recognition

1 Introduction

Human-centered artificial intelligence has become an important direction in the development of reliable, trustworthy, and privacy-aware intelligent systems [1, 2]. In sports applications, visual analysis is increasingly used for sport recognition, athlete monitoring, training analysis, performance assessment, and coaching support [3, 4]. Sports images collected during training, competition, and daily exercise contain rich information about athlete actions, body configurations, equipment interactions, and sport categories, providing an important basis for intelligent sports analysis. However, such data are often distributed across schools, clubs, teams, venues, and sports organizations. Directly centralizing them for model training may raise concerns regarding athlete privacy, institutional data ownership, and data governance. Therefore, sports visual intelligence requires learning mechanisms that can exploit distributed data while avoiding direct exposure of sensitive raw images [5, 6].

Federated Learning (FL) provides a promising solution by allowing multiple clients to collaboratively train a global model while keeping their raw data locally stored [7, 8]. This paradigm is well suited to sports applications in which different organizations may benefit from shared visual knowledge without directly exchanging athlete data. However, keeping raw data local does not provide rigorous privacy protection by itself. Previous studies have shown that transmitted gradients and model updates may still reveal sensitive training information through model inversion, membership inference, attribute inference, and data reconstruction attacks [9–12]. Differentially Private Federated Learning (DP-FL) therefore introduces gradient clipping and calibrated random noise into federated optimization to provide formal privacy guarantees [13–16].

Despite its privacy advantages, directly applying conventional DP-FL to sports image recognition remains challenging because distributed sports data often exhibit strong visual heterogeneity. Images of the same sport collected from different clients may share similar athlete poses, motion patterns, and human–equipment interactions, while differing considerably in athlete appearance, uniforms, backgrounds, illumination, viewpoints, and scene composition. The former represents transferable sport semantics that are useful across clients, whereas the latter often reflects client-specific visual context. An effective federated sports model should therefore preserve knowledge that characterizes the sport itself while reducing excessive dependence on visual patterns that are specific to individual clients.

Existing FL and DP-FL methods generally process local representations and model updates as a whole, without explicitly considering that different learned components may exhibit substantially different cross-client shareability. As a result, transferable sport semantics and client-specific contextual information can become entangled

during local optimization. Under heterogeneous data distributions, persistent differences in athlete appearance, uniforms, backgrounds, illumination, and viewpoints may dominate local representations and increase disagreement among client updates. When differential privacy is subsequently applied, clipping and noise perturbation are imposed uniformly on these mixed updates. Low-shareability contextual variations may therefore consume part of the limited clipping range, while transferable and discriminative sport information is unnecessarily distorted. This interaction between visual heterogeneity and uniform private processing can lead to substantial utility degradation in privacy-preserving federated sports recognition.

Existing DP-FL studies mainly alleviate privacy-induced utility degradation through improved optimization stability, clipping, sparsification, or perturbation efficiency. Sharpness-Aware Minimization (SAM), for example, has been incorporated into private federated optimization to encourage flatter solutions and improve robustness to noisy updates [17, 18]. Perturbation and sparsification strategies have also been explored to reduce the influence of high-dimensional privacy noise [19]. Gradient-Aware Minimization (GAM) further characterizes model stability using first-order flatness and provides a more sensitive description of the local optimization landscape [20]. Although these approaches improve private optimization, they generally treat the entire client update uniformly and do not explicitly distinguish globally transferable task knowledge from predominantly local visual context. This leaves an important question unresolved: *what knowledge should be shared across heterogeneous sports clients under privacy constraints?*

To address this problem, we propose **FedSS**, a shareability-aware privacy-preserving federated learning framework for sports image recognition. Instead of treating all learned components as equally suitable for federation, FedSS explicitly estimates their cross-client shareability and promotes globally transferable sport semantics while retaining client-specific visual context locally. At its core, a **Sport Semantic Shareability Module (SSSM)** evaluates stage-wise model components according to sport discriminativeness, cross-client invariance, and robustness to privacy perturbation. The resulting shareability scores guide private federated optimization by emphasizing highly shareable sport-semantic components before clipping and Gaussian perturbation, while suppressing low-shareability contextual variations. First-order gradient stability is further exploited to improve the robustness of transferable updates under privacy perturbation. In this way, FedSS reduces interference from client-specific visual context and preserves more transferable sport-semantic information within the limited clipping range, thereby jointly mitigating visual heterogeneity and privacy-induced utility degradation.

Our main contributions are summarized as follows:

- We propose **FedSS**, a shareability-aware privacy-preserving federated learning framework for sports image recognition under heterogeneous data distributions. FedSS explicitly distinguishes globally transferable sport semantics from client-specific visual context, enabling more effective knowledge sharing under privacy constraints.

- We introduce a **Sport Semantic Shareability Module** that adaptively evaluates stage-wise shareability according to sport discriminativeness, cross-client invariance, and robustness to privacy perturbation. The estimated shareability guides private federated optimization so that globally useful sport knowledge receives greater emphasis before clipping and perturbation, while low-shareability contextual information is retained locally.
- We conduct comprehensive experiments on sports image recognition benchmarks under varying client scales, participation rates, privacy budgets, and visual heterogeneity settings. Extensive results and diagnostic analyses demonstrate improved privacy–utility trade-offs, reduced clipping distortion, and stronger cross-client sport-semantic consistency compared with representative DP-FL baselines.

2 Related Work

2.1 Federated Sports Recognition under Visual Heterogeneity

Artificial intelligence has been widely applied to sports image recognition, athlete monitoring, training analysis, and performance assessment [3, 4, 21]. In practical applications, sports images are often distributed across teams, schools, clubs, and venues, making federated learning a suitable paradigm for collaborative model training without centralized collection of raw data. Existing federated learning studies mainly investigate statistical heterogeneity such as label-distribution skew [8, 22–24], while visual applications may additionally exhibit substantial feature-level shifts. FedBN, for example, preserves local batch-normalization statistics to mitigate feature distribution differences across clients [25]. In sports recognition, even clients containing the same sport categories may differ substantially in athlete appearance, uniforms, backgrounds, illumination, viewpoints, and scene composition. Such variations can produce inconsistent local representations, whereas sport-related patterns such as actions, body configurations, and human–equipment interactions are generally more transferable across clients. Therefore, distinguishing transferable sport semantics from client-specific visual context is important for federated sports recognition under heterogeneous data distributions.

2.2 Differentially Private Federated Learning

Federated learning enables distributed clients to collaboratively optimize a global model while retaining raw data locally [7, 8], but shared gradients and model updates may still expose sensitive information through model inversion, membership inference, attribute inference, and gradient reconstruction attacks [9–12, 26–29]. Differential privacy has therefore been widely incorporated into federated learning to provide formal privacy guarantees [13–16, 30–32]. A central challenge of DP-FL is the utility degradation caused by update clipping and random perturbation. Existing studies improve the privacy–utility trade-off through sparsification, adaptive optimization, adaptive clipping, sensitivity estimation, and flatness-aware training [17–19, 33–38]. These approaches mainly focus on how updates are optimized, clipped, or perturbed, while

comparatively less attention has been paid to whether different learned components have different values for cross-client sharing.

2.3 Shared and Client-Specific Knowledge in Federated Learning

The separation of globally shared and client-specific knowledge has been extensively studied in personalized federated learning. FedRep learns a common representation while maintaining client-specific prediction heads [39], Fed-RoD decouples generic and personalized prediction objectives through global and local predictors [40], and CD²-pFed performs channel-wise decomposition of shared and private model components [41]. FedBN similarly preserves client-specific batch-normalization statistics locally to accommodate feature distribution shifts [25]. These approaches demonstrate that different forms of client knowledge do not necessarily need to be globally aggregated. However, the shared-private separation is typically determined by predefined architectural components, layers, channels, prediction heads, or feature statistics, and is mainly designed for personalization. How to determine whether learned information is simultaneously transferable across heterogeneous clients and suitable for sharing under differential privacy remains less explored.

3 Preliminaries

3.1 Federated Sports Image Classification

We consider a federated sports image classification system consisting of M clients. Each client corresponds to a local sports data owner, such as a school, sports club, training center, or edge device. The local dataset of client i is denoted by

$$\mathcal{D}_i = (x_{ij}, y_{ij})_{j=1}^{n_i}, \quad (1)$$

where x_{ij} denotes a sports image and y_{ij} is its sport-category label. Raw images remain stored locally throughout federated training.

The clients collaboratively optimize a global classifier $f_{\mathbf{w}}$ with parameters $\mathbf{w}$. The standard federated objective is

$$\min_{\mathbf{w}} F(\mathbf{w}) = \sum_{i=1}^M \frac{n_i}{n} F_i(\mathbf{w}), \quad (2)$$

where $n = \sum_{i=1}^M n_i$ and

$$F_i(\mathbf{w}) = \frac{1}{n_i} \sum_{j=1}^{n_i} \ell(f_{\mathbf{w}}(x_{ij}), y_{ij}). \quad (3)$$

Unlike conventional IID learning, federated sports data exhibit both label heterogeneity and substantial visual-context heterogeneity caused by different athletes, uniforms, venues, backgrounds, viewpoints, illumination conditions, and acquisition devices.

3.2 Client-level Differential Privacy

Although federated learning avoids direct transmission of raw sports images, model updates can still reveal information about local training data. We therefore consider client-level differential privacy, where the contribution of each participating client is bounded before being transmitted to the server.

Given a local update Δ_i^t at communication round t , L_2 clipping produces

$$\hat{\Delta}_i^t = \Delta_i^t \cdot \min \left(1, \frac{C}{\|\Delta_i^t\|_2} \right), \quad (4)$$

where C is the clipping threshold. Gaussian noise is subsequently added as

$$\tilde{\Delta}_i^t = \hat{\Delta}_i^t + \mathcal{N}(0, \sigma^2 C^2 \mathbf{I}), \quad (5)$$

where σ is the noise multiplier. Under a specified client sampling rate and privacy accountant, the resulting mechanism provides client-level (ϵ, δ) -DP.

3.3 Gradient-Aware Minimization

Gradient-Aware Minimization (GAM) [20] characterizes local model stability through first-order flatness. Given a local loss $L_i(\mathbf{w})$, GAM optimizes

$$\min_{\mathbf{w}} L_i(\mathbf{w}) + \lambda \max_{\|\mathbf{s}\|_2 \leq \rho} \|\nabla L_i(\mathbf{w} + \mathbf{s})\|_2, \quad (6)$$

where ρ denotes the neighborhood radius and λ controls the first-order flatness regularization. Compared with optimization based only on the loss value within a perturbation neighborhood, first-order flatness directly characterizes the stability of gradient directions. In FedSS, GAM is used as a local optimization primitive and also provides gradient-stability information for estimating the robustness of candidate shared knowledge.

4 Method

4.1 Overview of FedSS

Figure 1 presents the overall framework of **FedSS**. Sports data distributed across clients contain both transferable sport semantics, such as actions, body configurations, and human–equipment interactions, and client-specific visual context, such as athlete appearance, uniforms, backgrounds, and viewpoints. Conventional DP-FL privatizes the complete local update without distinguishing these two types of information, which may cause client-specific variations to consume part of the limited clipping range.




Fig. 1 Overview of FedSS. The server broadcasts the global model and privacy-protected stage-wise EMA references to participating clients. After GAM-based local optimization, the Sport Semantic Shareability Module estimates stage-wise shareability from sport discriminativeness, cross-client invariance, and privacy robustness, and softly separates the candidate update into transferable sport semantics and client-specific context. The context residual is retained locally, whereas only the transferable component is clipped, perturbed with Gaussian noise, and uploaded for private server aggregation. The aggregated update is used to construct the next global model and refresh the stage-wise EMA references.

FedSS instead organizes federation according to the *shareability* of learned knowledge. At each communication round, a selected client performs local optimization and obtains a candidate update. The **Sport Semantic Shareability Module (SSSM)** evaluates each network stage from three perspectives: sport discriminativeness, cross-client invariance, and privacy robustness. The resulting shareability scores softly separate the candidate update into a transferable sport-semantic component and a client-specific context component. Only the transferable component is privatized and uploaded, while the context component is retained locally. The server then aggregates the private sport-semantic updates to update the global model.

4.2 Sport Semantic Shareability Module

FedSS estimates shareability at the network-stage level to avoid the instability and computational cost of parameter-wise evaluation. For the ResNet-18 backbone, the network is divided into five stages,

$$\mathcal{B} = \{B_1, B_2, B_3, B_4, B_5\}, \quad (7)$$

where B_1 contains the stem and Layer1, B_2 – B_4 correspond to Layer2–Layer4, and B_5 denotes the classifier. The candidate update of client i at round t is correspondingly partitioned as

$$\Delta_i^t = [\Delta_{i,1}^t, \Delta_{i,2}^t, \dots, \Delta_{i,5}^t]. \quad (8)$$

A shareability coefficient $q_{i,b}^t \in [0, 1]$ is estimated for each stage B_b and updated dynamically across clients and communication rounds.

4.2.1 Sport Discriminativeness

Sport discriminativeness measures how strongly a network stage contributes to correct sport recognition. After local training, client i samples a calibration mini-batch $\mathcal{B}_i$ and computes the gradient–activation relevance

$$s_{i,b}^D = \frac{1}{|\mathcal{B}_i|} \sum_{(x,y) \in \mathcal{B}_i} \left\| \mathbf{a}_{i,b}(x) \odot \frac{\partial f_y(x)}{\partial \mathbf{a}_{i,b}(x)} \right\|_1, \quad (9)$$

where $\mathbf{a}_{i,b}(x)$ is the activation of stage B_b and $f_y(x)$ is the logit of the ground-truth category. The normalized discriminativeness score is

$$D_{i,b}^t = \frac{s_{i,b}^D}{\max_k s_{i,k}^D + \epsilon}. \quad (10)$$

A larger $D_{i,b}^t$ indicates a stronger contribution to sport recognition.

4.2.2 Cross-client Invariance

Local discriminativeness alone cannot determine whether information is suitable for global sharing. FedSS therefore measures the consistency between the current stage update and knowledge accumulated across clients. The server maintains an exponential moving average of previously released private updates:

$$\mathbf{m}_b^{t-1} = \lambda_m \mathbf{m}_b^{t-2} + (1 - \lambda_m) \overline{\Delta}_b^{t-1}, \quad (11)$$

where λ_m is the smoothing coefficient. The cross-client invariance score is

$$I_{i,b}^t = \frac{1}{2} \left(1 + \frac{\langle \Delta_{i,b}^t, \mathbf{m}_b^{t-1} \rangle}{\|\Delta_{i,b}^t\|_2 \|\mathbf{m}_b^{t-1}\|_2 + \epsilon} \right). \quad (12)$$

A larger value indicates stronger consistency with cross-client knowledge. For the first communication round, we set $I_{i,b}^0 = 0.5$.

4.2.3 Privacy Robustness

FedSS further evaluates whether a stage remains stable and informative under privacy perturbation. Let $\mathbf{g}_{i,b}^t$ and $\mathbf{g}_{i,b}^{t,+}$ denote the gradients before and after the GAM

perturbation, respectively. Their relative first-order variation is

$$\xi_{i,b}^t = \frac{\|\mathbf{g}_{i,b}^{t,+} - \mathbf{g}_{i,b}^t\|_2}{\|\mathbf{g}_{i,b}^t\|_2 + \epsilon}, \quad (13)$$

and the corresponding stability is

$$F_{i,b}^t = \exp(-\xi_{i,b}^t). \quad (14)$$

The relative signal strength with respect to Gaussian perturbation is approximated as

$$N_{i,b}^t = \frac{\|\Delta_{i,b}^t\|_2}{\|\Delta_{i,b}^t\|_2 + \sigma C \sqrt{d_b} + \epsilon}, \quad (15)$$

where d_b is the number of trainable parameters in stage B_b , C is the clipping threshold, and σ is the noise multiplier. The privacy-robustness score is then

$$R_{i,b}^t = F_{i,b}^t N_{i,b}^t. \quad (16)$$

4.2.4 Stage-wise Shareability Estimation

The three criteria are combined with equal weights:

$$S_{i,b}^t = \frac{D_{i,b}^t + I_{i,b}^t + R_{i,b}^t}{3}. \quad (17)$$

The final shareability coefficient is obtained using a soft threshold:

$$q_{i,b}^t = \text{Sigmoid} \left(\frac{S_{i,b}^t - \tau}{T} \right), \quad (18)$$

where τ controls the threshold and T controls the softness of the assignment.

For compact notation, the stage-wise shareability operator is defined as

$$\mathbf{Q}_i^t = \text{Diag} (q_{i,1}^t \mathbf{I}_{d_1}, \dots, q_{i,5}^t \mathbf{I}_{d_5}). \quad (19)$$

The coefficients are computed locally and are never directly transmitted to the server.

4.3 Shareability-Guided Semantic-Context Decoupling

FedSS uses the estimated shareability to softly decompose each candidate update into a transferable sport-semantic component and a client-specific context component:

$$\Delta_{i,\text{sport}}^t = \mathbf{Q}_i^t \Delta_i^t, \quad \Delta_{i,\text{context}}^t = (\mathbf{I} - \mathbf{Q}_i^t) \Delta_i^t. \quad (20)$$

Because $q_{i,b}^t$ is continuous, each network stage can contribute to both components rather than being assigned exclusively to global or local learning.

Each client maintains a local residual to preserve client-specific information:

$$\mathbf{r}_i^{t+1} = \mu \mathbf{r}_i^t + \Delta_{i,\text{context}}^t, \quad (21)$$

where μ controls context retention. When client i participates again, its local model is initialized as

$$\mathbf{w}_i^{t,0} = \mathbf{w}^t + \mathbf{r}_i^t. \quad (22)$$

The residual is retained locally and is never communicated.

4.4 Shareability-Guided Private Optimization

Starting from $\mathbf{w}_i^{t,0}$, each selected client performs E epochs of GAM-based local optimization and obtains

$$\Delta_i^t = \mathbf{w}_i^{t,E} - \mathbf{w}_i^{t,0}. \quad (23)$$

After semantic-context decoupling, client-level differential privacy is applied only to $\Delta_{i,\text{sport}}^t$. The transferable update is first clipped:

$$\hat{\Delta}_{i,\text{sport}}^t = \Delta_{i,\text{sport}}^t \min \left(1, \frac{C}{\|\Delta_{i,\text{sport}}^t\|_2} \right), \quad (24)$$

and Gaussian noise is then added:

$$\tilde{\Delta}_{i,\text{sport}}^t = \hat{\Delta}_{i,\text{sport}}^t + \mathbf{z}_i^t, \quad \mathbf{z}_i^t \sim \mathcal{N}(\mathbf{0}, \sigma^2 C^2 \mathbf{I}). \quad (25)$$

Only the privatized sport-semantic update is transmitted. Since the released update is clipped to the same norm bound before the Gaussian mechanism, FedSS follows the same client-level privacy accounting procedure as standard DP-FL. Filtering low-shareability components before clipping reduces their occupation of the clipping range and preserves more transferable signal under the same privacy configuration.

4.5 Server Aggregation

Given the participating client set $\mathcal{S}_t$, the server computes the aggregated private update

$$\bar{\Delta}^t = \sum_{i \in \mathcal{S}_t} \frac{n_i}{\sum_{j \in \mathcal{S}_t} n_j} \tilde{\Delta}_{i,\text{sport}}^t, \quad (26)$$

and updates the global model as

$$\mathbf{w}^{t+1} = \mathbf{w}^t + \bar{\Delta}^t. \quad (27)$$

The stage-wise reference used in the next communication round is updated by

$$\mathbf{m}_b^t = \lambda_m \mathbf{m}_b^{t-1} + (1 - \lambda_m) \bar{\Delta}_b^t, \quad b = 1, \dots, 5. \quad (28)$$

Since this reference is constructed entirely from already privatized updates, its reuse is post-processing of DP outputs.

Algorithm 1: FedSS

Input : Client datasets $\{\mathcal{D}_i\}_{i=1}^M$; rounds R ; local epochs E ; DP parameters (C, σ) ; GAM parameters; shareability parameters (τ, T) ; EMA coefficient λ_m ; context coefficient μ ; initial model $\mathbf{w}^0$

Output: Global model $\mathbf{w}^R$

- 1 Partition the backbone into stages via Eq. (7);
Initialize local residuals $\mathbf{r}_i^0 = \mathbf{0}$ and stage references $\mathbf{m}_b^{-1} = \mathbf{0}$;
- 2 **for** $t = 0$ **to** $R - 1$ **do**
- 3 Sample clients $\mathcal{S}_t$ and broadcast $\mathbf{w}^t$ and stage references;
- 4 **for** *each* $i \in \mathcal{S}_t$ *in parallel* **do**
- 5 Initialize the local model via Eq. (22);
- 6 Perform GAM-based local training for E epochs and obtain the candidate update via Eq. (23);
- 7 Partition the update via Eq. (8);
- 8 Compute sport discriminativeness via Eqs. (9)–(10);
- 9 Compute cross-client invariance via Eq. (12);
- 10 Compute privacy robustness via Eqs. (13)–(16);
- 11 Estimate stage-wise shareability via Eqs. (17)–(19);
- 12 Decompose the candidate update via Eq. (20);
- 13 Update the local residual via Eq. (21);
- 14 Privatize the transferable update via Eqs. (24)–(25);
- 15 Upload $\tilde{\Delta}_{i,\text{sport}}^t$;
- 16 **end**
- 17 Aggregate private updates via Eq. (26);
Update the global model via Eq. (27);
Update stage-wise references via Eq. (28);
- 18 **end**
- 19 **return** $\mathbf{w}^R$;

The complete training procedure is summarized in Algorithm 1.

5 Theoretical Analysis

We analyze FedSS from four complementary perspectives. First, we show that the data-dependent stage-wise shareability estimation, the cross-round local context residual, and the private EMA reference do not introduce an additional privacy channel as long as every released client update is independently bounded and privatized. Second, we establish that the stage-wise soft shareability operator is non-expansive and therefore

cannot increase L_2 clipping distortion. Third, we characterize the error of extracting transferable sport semantics from an update containing both sport-semantic and client-specific components. Finally, we derive an aggregation-level error bound that jointly accounts for semantic suppression, context leakage, clipping distortion, and Gaussian perturbation. Together, these results explain why FedSS can improve private federated utility without weakening the underlying client-level differential privacy guarantee.

Under the stage-wise implementation of FedSS, the shareability operator of client i at round t is

$$\mathbf{Q}_i^t = \text{Diag} (q_{i,1}^t \mathbf{I}_{d_1}, \dots, q_{i,5}^t \mathbf{I}_{d_5}), \quad (29)$$

where $q_{i,b}^t \in [0, 1]$ is the dynamically estimated shareability coefficient of stage B_b . Hence,

$$\mathbf{0} \preceq \mathbf{Q}_i^t \preceq \mathbf{I}, \quad (30)$$

and the same scalar coefficient is applied to all parameters belonging to a given network stage.

5.1 Privacy Guarantee

At round t , let $\mathcal{H}_{t-1}$ denote the public history available before client computation, including previously released private updates, global models, and the stage-wise EMA references $\{\mathbf{m}_b^{t-1}\}_{b=1}^5$. The local context residual $\mathbf{r}_i^t$, in contrast, remains private to client i .

Conditioned on $\mathcal{H}_{t-1}$, client i produces a candidate update Δ_i^t and a data-dependent stage-wise operator $\mathbf{Q}_i^t$. The transferable update is

$$\mathbf{u}_i^t = \mathbf{Q}_i^t \Delta_i^t, \quad (31)$$

which is clipped as

$$\hat{\mathbf{u}}_i^t = \text{Clip}_C (\mathbf{u}_i^t). \quad (32)$$

The released update is

$$\tilde{\mathbf{u}}_i^t = \hat{\mathbf{u}}_i^t + \mathbf{z}_i^t, \quad (33)$$

where the Gaussian perturbation is calibrated according to the clipping bound and the adopted client-level adjacency definition.

Proposition 1. *Conditioned on any public history $\mathcal{H}_{t-1}$, suppose that the stage-wise shareability coefficients and local context residual are never directly released, and that the only client-dependent quantity transmitted at round t is $\tilde{\mathbf{u}}_i^t$. Then the round- t mechanism can use the same client-level Gaussian-mechanism privacy calibration as standard DP-FL under the same clipping bound, client-sampling scheme, and adjacency definition. Adaptive composition over communication rounds therefore yields the same privacy-accounting procedure used by the underlying client-level DP-FL mechanism.*

Proof. For a fixed public history $\mathcal{H}_{t-1}$, define the deterministic pre-noise mechanism

$$h_t(\mathcal{D}_i) = \text{Clip}_C (\mathbf{Q}_i^t(\mathcal{D}_i) \Delta_i^t(\mathcal{D}_i)). \quad (34)$$

Although both the candidate update and the shareability coefficients may depend arbitrarily on the private local dataset and the private local state, clipping guarantees

$$\|h_t(\mathcal{D}_i)\|_2 \leq C. \quad (35)$$

Hence, under add/remove client adjacency with the absent-client contribution represented by zero, the sensitivity is bounded by C . Under replacement adjacency, the distance between two bounded contributions is at most $2C$. The corresponding Gaussian mechanism can therefore be calibrated using exactly the same sensitivity convention as the baseline client-level DP-FL algorithm.

The stage-wise scores $\{q_{i,b}^t\}_{b=1}^5$ and the local residual $\mathbf{r}_i^t$ remain internal variables and are never directly released. Although $\mathbf{r}_i^t$ affects future local initializations, every future client-dependent release is again passed through clipping and Gaussian perturbation. Differential privacy is closed under such adaptive composition of conditionally private mechanisms.

Furthermore, the EMA references $\{\mathbf{m}_b^{t-1}\}_{b=1}^5$ are constructed exclusively from previously released private updates. Their construction, broadcast, and reuse therefore constitute post-processing of differentially private outputs and incur no additional privacy loss. Applying the adopted subsampling and composition accountant over all communication rounds yields the final (ε, δ) guarantee. $\square$

This result is particularly important for FedSS because both the stage-wise shareability computation and the cross-round local memory are data dependent. Proposition 1 shows that these internal computations do not require separate privacy accounting as long as the externally released client contribution is always bounded and privatized.

5.2 Non-expansiveness and Clipping Distortion

A major source of utility loss in client-level DP-FL is clipping. For an arbitrary update $\mathbf{v}$, define its clipping distortion under threshold C as

$$\mathcal{E}_C(\mathbf{v}) = \|\mathbf{v} - \text{Clip}_C(\mathbf{v})\|_2. \quad (36)$$

For standard L_2 clipping,

$$\mathcal{E}_C(\mathbf{v}) = (\|\mathbf{v}\|_2 - C)_+, \quad (37)$$

where $(a)_+ = \max(a, 0)$.

The stage-wise structure of FedSS gives the following contraction property.

Proposition 2. *For any candidate update Δ_i and stage-wise shareability operator $\mathbf{Q}_i$ with $q_{i,b} \in [0, 1]$,*

$$\|\mathbf{Q}_i \Delta_i\|_2 \leq \|\Delta_i\|_2. \quad (38)$$

Consequently,

$$\mathcal{E}_C(\mathbf{Q}_i \Delta_i) \leq \mathcal{E}_C(\Delta_i). \quad (39)$$

Proof. Decomposing Δ_i according to the five network stages gives

$$\Delta_i = [\Delta_{i,1}, \dots, \Delta_{i,5}]. \quad (40)$$

By the definition of $\mathbf{Q}_i$,

$$\begin{aligned}\|\mathbf{Q}_i \Delta_i\|_2^2 &= \sum_{b=1}^5 q_{i,b}^2 \|\Delta_{i,b}\|_2^2 \\ &\leq \sum_{b=1}^5 \|\Delta_{i,b}\|_2^2 = \|\Delta_i\|_2^2,\end{aligned}\tag{41}$$

because $0 \leq q_{i,b} \leq 1$. Therefore, $\|\mathbf{Q}_i \Delta_i\|_2 \leq \|\Delta_i\|_2$. Since $(x - C)_+$ is monotonically non-decreasing in x , the clipping distortion inequality follows immediately. $\square$

Corollary 1. *If $\|\Delta_i\|_2 > C$ and at least one stage with a non-zero update is strictly attenuated, i.e., $q_{i,b} < 1$, then $\|\mathbf{Q}_i \Delta_i\|_2 < \|\Delta_i\|_2$. The clipping distortion is therefore strictly reduced whenever the filtered update remains above C , and becomes zero if its norm falls below C .*

Proposition 2 does not claim that reducing the update norm alone is sufficient for better learning. Rather, it establishes a guaranteed optimization-side benefit of the SSSM: stage-wise filtering cannot make clipping more severe. The semantic value of the filtering decision is characterized next.

5.3 Stage-wise Semantic–Context Separation

To characterize the semantic effect of stage-wise shareability, we conceptually decompose the candidate update of client i as

$$\Delta_i = \Delta_i^s + \Delta_i^c,\tag{42}$$

where Δ_i^s denotes the ideal transferable sport-semantic component and Δ_i^c denotes client-specific contextual variation. Both components follow the same five-stage partition:

$$\Delta_i^s = [\Delta_{i,1}^s, \dots, \Delta_{i,5}^s], \quad \Delta_i^c = [\Delta_{i,1}^c, \dots, \Delta_{i,5}^c].\tag{43}$$

This decomposition is used only for analysis and is not assumed to be available to FedSS during training.

FedSS selects

$$\mathbf{u}_i = \mathbf{Q}_i \Delta_i.\tag{44}$$

Its deviation from the ideal semantic update satisfies

$$\begin{aligned}\mathbf{u}_i - \Delta_i^s &= \mathbf{Q}_i (\Delta_i^s + \Delta_i^c) - \Delta_i^s \\ &= -(\mathbf{I} - \mathbf{Q}_i) \Delta_i^s + \mathbf{Q}_i \Delta_i^c.\end{aligned}\tag{45}$$

The first term corresponds to transferable sport semantics suppressed by the stage-wise gate, while the second corresponds to client-specific context that remains in the shared update.

Proposition 3. *For any stage-wise shareability operator $\mathbf{Q}_i$,*

$$\|\mathbf{u}_i - \Delta_i^s\|_2 \leq \|(\mathbf{I} - \mathbf{Q}_i) \Delta_i^s\|_2 + \|\mathbf{Q}_i \Delta_i^c\|_2.\tag{46}$$

Proof. The result follows directly from the decomposition above and the triangle inequality:

$$\begin{aligned}\|\mathbf{u}_i - \Delta_i^s\|_2 &= \|-(\mathbf{I} - \mathbf{Q}_i)\Delta_i^s + \mathbf{Q}_i\Delta_i^c\|_2 \\ &\leq \|(\mathbf{I} - \mathbf{Q}_i)\Delta_i^s\|_2 + \|\mathbf{Q}_i\Delta_i^c\|_2.\end{aligned}\quad (47)$$

□

For non-zero semantic and contextual components, define the semantic suppression ratio

$$\kappa_i^s = \frac{\|(\mathbf{I} - \mathbf{Q}_i)\Delta_i^s\|_2}{\|\Delta_i^s\|_2}, \quad (48)$$

and the context leakage ratio

$$\kappa_i^c = \frac{\|\mathbf{Q}_i\Delta_i^c\|_2}{\|\Delta_i^c\|_2}. \quad (49)$$

If either denominator is zero, its corresponding ratio is defined as zero. Since $\mathbf{0} \preceq \mathbf{Q}_i \preceq \mathbf{I}$, both ratios lie in $[0, 1]$. Proposition 3 therefore becomes

$$\|\mathbf{u}_i - \Delta_i^s\|_2 \leq \kappa_i^s \|\Delta_i^s\|_2 + \kappa_i^c \|\Delta_i^c\|_2. \quad (50)$$

This expression exposes the two competing objectives of semantic shareability. A small κ_i^s means that useful sport knowledge is preserved, whereas a small κ_i^c means that client-specific context is suppressed.

The same interpretation extends directly to server aggregation. Let

$$\omega_i = \frac{n_i}{\sum_{j \in \mathcal{S}_t} n_j} \quad (51)$$

denote the aggregation weight of client i . Then the contextual component entering the server before clipping satisfies

$$\left\| \sum_{i \in \mathcal{S}_t} \omega_i \mathbf{Q}_i \Delta_i^c \right\|_2 \leq \sum_{i \in \mathcal{S}_t} \omega_i \kappa_i^c \|\Delta_i^c\|_2. \quad (52)$$

Because $\kappa_i^c \leq 1$, FedSS provides a no-larger worst-case upper bound on contextual contamination than sharing the complete client updates. The reduction is strict whenever context-dominated stages are assigned coefficients below one.

The three components of the SSSM are designed to reduce these two error sources from complementary perspectives. Sport discriminativeness encourages retention of stages that contribute to correct sport prediction, cross-client invariance penalizes update directions that are inconsistent with the privacy-protected cross-client history, and privacy robustness favors stages whose signals remain stable under local perturbation and DP noise. Importantly, Proposition 3 does not assume that FedSS observes the latent semantic/context decomposition; it only characterizes the desirable behavior induced when the estimated shareability coefficients align with that decomposition.

5.4 Error of Private Server Aggregation

We finally connect semantic separation, clipping, and Gaussian perturbation at the server level. For each participating client, define

$$\mathbf{u}_i = \mathbf{Q}_i \Delta_i, \quad \hat{\mathbf{u}}_i = \text{Clip}_C(\mathbf{u}_i), \quad (53)$$

and

$$\tilde{\mathbf{u}}_i = \hat{\mathbf{u}}_i + \mathbf{z}_i, \quad (54)$$

where

$$\mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, \nu^2 \mathbf{I}) \quad (55)$$

and the noises are independent across participating clients.

The aggregated private update is

$$\tilde{\mathbf{g}} = \sum_{i \in \mathcal{S}_t} \omega_i \tilde{\mathbf{u}}_i, \quad (56)$$

whereas the ideal semantic aggregate is

$$\mathbf{g}^s = \sum_{i \in \mathcal{S}_t} \omega_i \Delta_i^s. \quad (57)$$

Proposition 4. *Let d denote the dimensionality of the transmitted model update. Then*

$$\begin{aligned} \mathbb{E} [\|\tilde{\mathbf{g}} - \mathbf{g}^s\|_2^2] &\leq \left[\sum_{i \in \mathcal{S}_t} \omega_i (\kappa_i^s \|\Delta_i^s\|_2 + \kappa_i^c \|\Delta_i^c\|_2 + \mathcal{E}_C(\mathbf{u}_i)) \right]^2 \\ &\quad + d\nu^2 \sum_{i \in \mathcal{S}_t} \omega_i^2. \end{aligned} \quad (58)$$

Proof. Write the aggregation error as

$$\begin{aligned} \tilde{\mathbf{g}} - \mathbf{g}^s &= \sum_{i \in \mathcal{S}_t} \omega_i (\hat{\mathbf{u}}_i - \Delta_i^s) \\ &\quad + \sum_{i \in \mathcal{S}_t} \omega_i \mathbf{z}_i. \end{aligned} \quad (59)$$

Because the Gaussian noises are zero-mean and independent of the deterministic clipped updates,

$$\begin{aligned} \mathbb{E} [\|\tilde{\mathbf{g}} - \mathbf{g}^s\|_2^2] &= \left\| \sum_{i \in \mathcal{S}_t} \omega_i (\hat{\mathbf{u}}_i - \Delta_i^s) \right\|_2^2 \\ &\quad + \mathbb{E} \left\| \sum_{i \in \mathcal{S}_t} \omega_i \mathbf{z}_i \right\|_2^2. \end{aligned} \quad (60)$$

For the deterministic term, the triangle inequality gives

$$\begin{aligned}\|\hat{\mathbf{u}}_i - \Delta_i^s\|_2 &\leq \|\hat{\mathbf{u}}_i - \mathbf{u}_i\|_2 + \|\mathbf{u}_i - \Delta_i^s\|_2 \\ &\leq \mathcal{E}_C(\mathbf{u}_i) + \kappa_i^s \|\Delta_i^s\|_2 + \kappa_i^c \|\Delta_i^c\|_2,\end{aligned}\tag{61}$$

where the second inequality follows from Proposition 3. Applying the triangle inequality once more after aggregation yields the squared first term in the proposition.

For the stochastic term, independence and isotropic Gaussian covariance give

$$\mathbb{E} \left\| \sum_{i \in \mathcal{S}_t} \omega_i \mathbf{z}_i \right\|_2^2 = d\nu^2 \sum_{i \in \mathcal{S}_t} \omega_i^2.\tag{62}$$

Combining the deterministic and stochastic terms proves the result. $\square$

Proposition 4 reveals four distinct sources of error in the private global update: semantic suppression, context leakage, clipping distortion, and the unavoidable Gaussian noise floor. Under a fixed client-level privacy configuration, the last term is determined by the DP mechanism and cannot be directly eliminated by FedSS. The role of FedSS is instead to reduce the first three terms before private aggregation. The stage-wise SSSM is designed to simultaneously preserve sport-discriminative knowledge and suppress client-specific context, while Proposition 2 guarantees that the resulting soft filtering cannot increase clipping distortion. The privacy-robustness criterion further biases sharing toward stages whose optimization signals are stable relative to DP perturbation. This decomposition therefore provides a direct theoretical explanation for the empirical improvements in cross-client semantic consistency, clipping behavior, and privacy-utility trade-off observed in our experiments.

6 Experiments

6.1 Experimental Setup

6.1.1 Datasets

We evaluate FedSS on two sports image recognition benchmarks, sport22 [42] and sport100 [43]. Sport22 contains 22 sports categories, whereas sport100 contains 100 categories and provides a more fine-grained recognition setting. The two datasets cover diverse athlete appearances, sports equipment, backgrounds, and acquisition conditions, making them suitable for studying privacy-preserving federated sports recognition. Following the federated setting, the training images are distributed across multiple clients, while the original test split of each dataset is retained for evaluating the global sports classifier. Raw sports images remain stored locally throughout federated training.

6.1.2 Federated Sports Heterogeneity

A central objective of our experiments is to distinguish conventional label heterogeneity from persistent client-specific visual heterogeneity. We therefore consider four

Table 1 Federated sports heterogeneity settings used in our experiments.

Setting	Label Distribution	Visual Context
IID	Balanced	Shared
L-Het	Non-IID ($\alpha = 0.3$)	Shared
C-Het	Balanced	Client-specific
M-Het	Non-IID ($\alpha = 0.3$)	Client-specific

data-distribution settings. IID approximately balances both sport categories and visual conditions across clients. L-Het introduces label heterogeneity using a Dirichlet partition with $\alpha = 0.3$ while keeping visual conditions shared. C-Het maintains approximately balanced class distributions but assigns each client a fixed visual-context profile consisting of illumination, color style, viewpoint, and camera-quality transformations. The parameters of each profile are sampled once and remain unchanged throughout training, thereby simulating persistent acquisition differences across sports data owners rather than random data augmentation. The same client-specific profile is also applied to the corresponding held-out test samples. M-Het combines the Dirichlet label partition with these persistent client-specific context profiles, simultaneously introducing semantic-distribution skew and visual-context shift. We further consider two system scales. In **S1**, the dataset is distributed across 100 clients and 10% of the clients participate in each communication round. In **S2**, the dataset is distributed across 500 clients and 2% participate in each round. Both settings involve approximately the same number of active clients per round, while S2 represents a more fragmented federation with substantially lower participation frequency for each individual client.

6.1.3 Baselines

We compare FedSS with several representative privacy-preserving federated learning methods, including DP-FedAvg [44], which applies client-level clipping and Gaussian perturbation to conventional FedAvg updates; DP-FedAvg-blur [17], which improves private optimization through gradient smoothing; Fed-SMP [19], which adopts sparse model perturbation for client-level differentially private federated learning; DP-FedSAM [18], which incorporates Sharpness-Aware Minimization into private federated optimization; ADDPNFL [33], an adaptive private federated optimizer using first- and second-order moment information; and DPSIR-Fed [38], which employs sparsity-inducing reparameterization and activation-pruned backpropagation to concentrate salient update signals and improve robustness to differential privacy noise.

6.1.4 Implementation Details

For both datasets, we adopt ResNet-18 as the backbone. The classification head contains 22 outputs for sport22 and 100 outputs for sport100. Batch Normalization layers

are replaced with Group Normalization to improve stability under heterogeneous federated training [45]. Unless otherwise specified, the total number of communication rounds is $R = 400$, the number of local epochs is $E = 20$, and the local batch size is 50. The initial learning rate is 0.1 and the weight decay is 5×10^{-4} . For GAM-based optimization, the perturbation radius is set to $\rho = 0.5$ by default. The clipping threshold is fixed to $C = 0.4$, and the default Gaussian noise multiplier is $\sigma = 0.95$. The failure probability is $\delta = 1/M$, where M denotes the total number of federated clients. The privacy budget ε is calculated using a Rényi Differential Privacy accountant according to the client sampling rate, communication rounds, noise multiplier, and δ . For the SSSM, the default shareability threshold is $\tau = 0.5$, the temperature is $T = 0.2$, and the three shareability terms are assigned equal initial weights unless otherwise specified. The local context-retention coefficient is set to $\mu = 0.75$. The model update is partitioned into parameter blocks corresponding to network stages and classifier parameters when computing shareability. All quantitative results are reported as mean $\pm$ standard deviation over three independent runs with different random seeds.

Unless otherwise specified, all main experiments use $\sigma = 0.95$, $C = 0.4$, and $\delta = 1/M$. Under S1 with a 10% client-sampling rate and $R = 400$ communication rounds, the RDP accountant gives a default privacy budget of approximately $\varepsilon = 12$. The experiments in Table 3 separately vary the privacy budget by adjusting the Gaussian noise multiplier while keeping the remaining federated configuration unchanged.

6.1.5 Evaluation Metrics

We evaluate recognition utility, privacy protection, representation quality, and update behavior using the following metrics.

Top-1 Accuracy. We report Top-1 classification accuracy on the original test set as the primary measure of recognition performance.

Privacy Level. Client-level privacy is measured by the privacy budget ε under a fixed $\delta = 1/M$, where M is the total number of clients. Smaller ε indicates stronger privacy protection.

Sport Semantic Consistency (SCS). SCS measures the cross-client consistency of representations belonging to the same sport category. For each valid class c , we compute client class prototypes $\mathbf{p}_{i,c}$ and average their pairwise cosine similarity:

$$\text{SCS} = \frac{1}{|\mathcal{C}'|} \sum_{c \in \mathcal{C}'} \frac{1}{|\mathcal{P}_c|} \sum_{(i,j) \in \mathcal{P}_c} \cos(\mathbf{p}_{i,c}, \mathbf{p}_{j,c}), \quad (63)$$

where $\mathcal{P}_c$ contains valid client pairs sharing class c . Higher SCS indicates stronger cross-client semantic consistency.

Class-Conditional Client-Context Predictability (C-CCP). C-CCP measures how much client-specific information remains in the shared representation after controlling for sport category. For each class c , a linear probe predicts client identity from frozen representations, and the balanced accuracy A_c is normalized by its chance

level:

$$\text{C-CCP} = \frac{100}{|\mathcal{C}'|} \sum_{c \in \mathcal{C}'} \frac{A_c - 1/K_c}{1 - 1/K_c}, \quad (64)$$

where K_c is the number of eligible clients for class c . Lower C-CCP indicates less client-specific information in the shared representation.

Pre-clipping Update Norm. We report the average L_2 norm of the update immediately before differential privacy clipping. For FedSS, we additionally compare the candidate-update norm before shareability filtering with the transferable-update norm after filtering.

Clipping Rate. The clipping rate is the proportion of participating clients whose pre-clipping update norm exceeds the clipping threshold C :

$$\text{ClipRate} = \frac{1}{|\mathcal{S}_t|} \sum_{i \in \mathcal{S}_t} \mathbb{I} [\|\mathbf{u}_i^t\|_2 > C]. \quad (65)$$

A lower value indicates that fewer updates require clipping.

Clipping Distortion. Clipping distortion measures the information removed by norm clipping:

$$\text{ClipDist} = \frac{1}{|\mathcal{S}_t|} \sum_{i \in \mathcal{S}_t} \|\mathbf{u}_i^t - \text{Clip}_C(\mathbf{u}_i^t)\|_2. \quad (66)$$

Lower ClipDist indicates that more of the original update is preserved before Gaussian perturbation.

Implementation and Reproducibility Details.

All experiments are implemented in Python 3.10 using PyTorch 2.1.2 and CUDA 12.1, and are conducted on a workstation equipped with an NVIDIA GeForce RTX 4090 GPU with 24 GB memory, an Intel Xeon CPU, and 128 GB system RAM, running Ubuntu 22.04. Unless otherwise specified, all reported results are averaged over three independent runs with random seeds 20, 21, and 22, and the corresponding standard deviations are reported. The same random seeds are used across different methods under each experimental configuration to ensure a consistent comparison. Dataset partitioning, client sampling, model initialization, and other stochastic operations are controlled by the corresponding random seed. We use ResNet-18 as the backbone and replace Batch Normalization with Group Normalization to reduce the influence of client-specific batch statistics under heterogeneous federated settings. All methods are implemented within the same federated learning codebase and evaluated using identical data partitions, client participation schedules, and privacy-accounting protocols whenever applicable. The implementation code and experiment configurations are available internally and will be made publicly available upon reasonable request and/or following publication to facilitate reproducibility.

6.2 Main Results under Structured Sports Heterogeneity

Table 2 compares the Top-1 accuracy under four heterogeneity settings using S1. FedSS achieves the best performance across all settings on both sport22 and sport100. Compared with the strongest baseline, DPSIR-Fed, FedSS provides the largest gains under

Table 2 Top-1 test accuracy (%) under different heterogeneity settings using S1. L-Het denotes label heterogeneity, C-Het denotes visual-context heterogeneity, and M-Het denotes their combination.

sport22				
Method	IID	L-Het	C-Het	M-Het
DP-FedAvg	46.19 $\pm$ 0.17	41.43 $\pm$ 0.31	40.16 $\pm$ 0.42	37.72 $\pm$ 0.48
DP-FedAvg-blur	47.54 $\pm$ 0.38	43.23 $\pm$ 0.53	42.08 $\pm$ 0.44	39.71 $\pm$ 0.51
Fed-SMP-top _k	49.54 $\pm$ 0.29	45.99 $\pm$ 0.67	44.67 $\pm$ 0.39	42.35 $\pm$ 0.46
DP-FedSAM	53.12 $\pm$ 0.56	48.98 $\pm$ 0.92	47.31 $\pm$ 0.53	44.86 $\pm$ 0.59
ADDPNFL	52.32 $\pm$ 0.32	47.98 $\pm$ 0.99	46.12 $\pm$ 0.48	43.57 $\pm$ 0.55
DPSIR-Fed	57.63 $\pm$ 0.32	52.76 $\pm$ 0.41	49.98 $\pm$ 0.38	47.82 $\pm$ 0.44
FedSS	59.84 $\pm$ 0.28	54.91 $\pm$ 0.35	55.76 $\pm$ 0.31	50.93 $\pm$ 0.37
sport100				
Method	IID	L-Het	C-Het	M-Het
DP-FedAvg	39.15 $\pm$ 0.37	35.33 $\pm$ 0.12	34.07 $\pm$ 0.35	31.86 $\pm$ 0.41
DP-FedAvg-blur	39.98 $\pm$ 0.26	36.58 $\pm$ 0.42	35.46 $\pm$ 0.38	33.14 $\pm$ 0.43
Fed-SMP-top _k	43.43 $\pm$ 0.28	38.95 $\pm$ 0.76	38.18 $\pm$ 0.44	35.02 $\pm$ 0.48
DP-FedSAM	47.56 $\pm$ 0.71	41.94 $\pm$ 0.28	41.03 $\pm$ 0.42	38.71 $\pm$ 0.46
ADDPNFL	45.15 $\pm$ 0.79	41.65 $\pm$ 0.46	39.84 $\pm$ 0.47	37.29 $\pm$ 0.51
DPSIR-Fed	52.13 $\pm$ 0.50	47.72 $\pm$ 0.17	45.76 $\pm$ 0.36	43.08 $\pm$ 0.39
FedSS	53.84 $\pm$ 0.34	50.81 $\pm$ 0.29	51.47 $\pm$ 0.33	45.72 $\pm$ 0.35

C-Het, improving accuracy by 5.78 percentage points on sport22 and 5.71 points on sport100. This result indicates that explicitly modeling cross-client shareability is particularly effective when clients share similar sport categories but differ substantially in visual context. Under the more challenging M-Het setting, FedSS still outperforms DPSIR-Fed by 3.11 and 2.64 percentage points on sport22 and sport100, respectively. Overall, the results show that shareability-aware federation provides consistent benefits under both statistical and visual heterogeneity by better preserving transferable sport-semantic information during private aggregation.

6.3 Privacy–Utility Trade-off

We further evaluate FedSS under different privacy budgets $\varepsilon \in 4, 6, 8, 10, 12$ using the challenging S1-M-Het setting, with $\varepsilon = 12$ as the default configuration. As shown in Table 3, the accuracy of all methods generally increases with a larger privacy budget, while FedSS consistently achieves the best performance on both datasets. Compared with DPSIR-Fed, FedSS provides gains of approximately 3.0–4.1 percentage points on sport22 and 2.3–2.8 points on sport100 across the evaluated privacy budgets. The advantage is particularly clear under the strongest privacy setting ($\varepsilon = 4$), where FedSS improves accuracy from 42.6746.81% on sport22 and from 36.79% to 39.63% on sport100. These results indicate that shareability-aware federation better preserves transferable sport-semantic information under stronger privacy perturbation, while maintaining consistent benefits across a broad range of privacy budgets.

Table 3 Top-1 accuracy (%) under different privacy budgets in the S1-M-Het setting. Results are reported as mean $\pm$ standard deviation over three independent runs. Smaller ϵ indicates stronger privacy.

sport22					
Method	$\epsilon = 4$	$\epsilon = 6$	$\epsilon = 8$	$\epsilon = 10$	$\epsilon = 12$
DP-FedAvg	32.94 $\pm$ 0.61	34.86 $\pm$ 0.55	36.12 $\pm$ 0.52	37.15 $\pm$ 0.49	37.52 $\pm$ 0.48
DP-FedAvg-blur	34.25 $\pm$ 0.64	36.55 $\pm$ 0.60	37.62 $\pm$ 0.57	38.84 $\pm$ 0.54	39.71 $\pm$ 0.51
Fed-SMP-top $_k$	36.90 $\pm$ 0.59	39.10 $\pm$ 0.55	40.55 $\pm$ 0.51	41.54 $\pm$ 0.48	42.35 $\pm$ 0.46
DP-FedSAM	38.71 $\pm$ 0.74	40.91 $\pm$ 0.69	42.55 $\pm$ 0.65	43.88 $\pm$ 0.62	44.86 $\pm$ 0.59
ADDPNFL	37.96 $\pm$ 0.69	40.14 $\pm$ 0.65	41.68 $\pm$ 0.61	42.82 $\pm$ 0.58	43.57 $\pm$ 0.55
DPSIR-Fed	42.67 $\pm$ 0.57	44.53 $\pm$ 0.53	45.82 $\pm$ 0.49	46.96 $\pm$ 0.46	47.92 $\pm$ 0.44
FedSS	46.81$\pm$0.49	48.25$\pm$0.45	49.31$\pm$0.42	50.08$\pm$0.40	50.93$\pm$0.37

sport100					
Method	$\epsilon = 4$	$\epsilon = 6$	$\epsilon = 8$	$\epsilon = 10$	$\epsilon = 12$
DP-FedAvg	26.81 $\pm$ 0.57	28.66 $\pm$ 0.53	29.77 $\pm$ 0.48	30.91 $\pm$ 0.44	31.66 $\pm$ 0.45
DP-FedAvg-blur	28.32 $\pm$ 0.55	29.65 $\pm$ 0.51	30.81 $\pm$ 0.48	32.02 $\pm$ 0.45	33.14 $\pm$ 0.43
Fed-SMP-top $_k$	29.67 $\pm$ 0.62	31.21 $\pm$ 0.57	32.76 $\pm$ 0.54	33.91 $\pm$ 0.51	35.02 $\pm$ 0.48
DP-FedSAM	32.67 $\pm$ 0.58	34.25 $\pm$ 0.54	35.86 $\pm$ 0.51	37.31 $\pm$ 0.49	38.71 $\pm$ 0.45
ADDPNFL	31.95 $\pm$ 0.63	33.44 $\pm$ 0.58	34.96 $\pm$ 0.55	36.28 $\pm$ 0.53	37.29 $\pm$ 0.51
DPSIR-Fed	36.79 $\pm$ 0.52	38.71 $\pm$ 0.48	40.20 $\pm$ 0.45	41.72 $\pm$ 0.42	43.08 $\pm$ 0.39
FedSS	39.63$\pm$0.47	41.04$\pm$0.44	42.68$\pm$0.41	44.19$\pm$0.38	45.72$\pm$0.35

Table 4 Ablation study of FedSS under S1-M-Het. D , I , and R denote sport discriminativeness, cross-client invariance, and privacy robustness, respectively. SCS and C-CCP are reported on sport22.

Variant	sport22 Acc. $\uparrow$	sport100 Acc. $\uparrow$	SCS $\uparrow$	C-CCP (%) $\downarrow$	ClipDist $\downarrow$
w/o Shareability Module	46.82	42.08	0.701	34.86	0.109
w/o Sport Discriminativeness (D)	48.47	43.56	0.736	23.41	0.068
w/o Cross-client Invariance (I)	47.91	43.02	0.712	30.74	0.066
w/o Privacy Robustness (R)	49.16	44.11	0.753	22.68	0.081
w/o Local Context Memory	49.52	44.63	0.768	19.85	0.064
w/o GAM	49.87	44.91	0.773	20.31	0.072
Full FedSS	50.93	45.72	0.782	18.63	0.061

6.4 Ablation Study

We conduct an ablation study under S1-M-Het to evaluate the contribution of each FedSS component. As shown in Table 4, removing the complete SSSM causes the largest accuracy drop, from 50.93% to 46.82% on sport22 and from 45.72% to 42.08% on sport100, confirming that semantic shareability is the main source of improvement. The three shareability factors provide complementary benefits: sport discriminativeness mainly preserves task-relevant information, cross-client invariance improves semantic consistency and suppresses client-specific context, and privacy robustness reduces clipping distortion. Removing the local context memory also degrades accuracy, showing that low-shareability information is better retained locally

Table 5 Sensitivity to the shareability threshold τ and local context-retention coefficient μ under S1-M-Het. Results are reported as mean $\pm$ standard deviation over three independent runs.

Parameter	Value	sport22 Acc. (%)	sport100 Acc. (%)
τ	0.2	48.91 $\pm$ 0.46	43.97 $\pm$ 0.43
	0.4	50.35 $\pm$ 0.41	45.16 $\pm$ 0.39
	0.5	50.93$\pm$0.37	45.72$\pm$0.35
	0.6	50.21 $\pm$ 0.40	45.08 $\pm$ 0.38
	0.8	47.84 $\pm$ 0.49	42.96 $\pm$ 0.45
μ	0.00	49.52 $\pm$ 0.44	44.63 $\pm$ 0.41
	0.25	50.04 $\pm$ 0.42	45.02 $\pm$ 0.39
	0.50	50.48 $\pm$ 0.39	45.39 $\pm$ 0.37
	0.75	50.93$\pm$0.37	45.72$\pm$0.35
	0.90	50.36 $\pm$ 0.41	45.21 $\pm$ 0.38

than discarded. GAM provides an additional but smaller improvement in accuracy and clipping stability. Overall, the results confirm that SSSM is the core component of FedSS, while local context retention and first-order optimization provide complementary gains.

6.5 Sensitivity to Shareability Hyperparameters

We analyze the sensitivity of FedSS to the shareability threshold τ and local context retention coefficient μ under S1-M-Het. As shown in Table 5, FedSS remains relatively stable over a broad range of both hyperparameters, with moderate values providing the best performance. For τ , performance first improves as the threshold increases, reaching the best results at $\tau = 0.5$, while overly small or large values weaken the balance between preserving transferable sport semantics and suppressing low-shareability context. A similar trend is observed for μ , where moderate context retention benefits local adaptation, but excessive retention slightly degrades performance. Overall, $\tau = 0.5$ and $\mu = 0.75$ provide the best results and are used as the default settings.

6.6 Scalability under Different Client Populations

As shown in Table 6, increasing the federation scale from S1 to S2 reduces the accuracy of all methods due to the larger client population and lower participation frequency. Nevertheless, FedSS consistently achieves the best performance on both datasets. Compared with DPSIR-Fed, FedSS maintains gains of about 3 percentage points on sport22 and 2.6 percentage points on sport100 under both S1 and S2. Although its absolute accuracy decreases as the federation scale increases, the relative advantage remains nearly unchanged. These results indicate that shareability-aware federation remains effective under sparse participation and larger-scale federated settings.

Table 6 Top-1 accuracy (%) under different client scales in the M-Het setting. S1 uses 100 clients with 10% participation, while S2 uses 500 clients with 2% participation. Results are reported as mean $\pm$ standard deviation over three independent runs.

Method	sport22		sport100	
	S1	S2	S1	S2
DP-FedAvg	37.72 $\pm$ 0.48	35.41 $\pm$ 0.65	31.86 $\pm$ 0.41	29.84 $\pm$ 0.49
DP-FedAvg-blur	39.71 $\pm$ 0.51	37.19 $\pm$ 0.58	33.14 $\pm$ 0.43	31.11 $\pm$ 0.51
Fed-SMP-top _k	42.35 $\pm$ 0.46	39.76 $\pm$ 0.54	35.02 $\pm$ 0.48	32.73 $\pm$ 0.52
DP-FedSAM	44.86 $\pm$ 0.59	42.18 $\pm$ 0.62	38.71 $\pm$ 0.46	36.42 $\pm$ 0.53
ADDPNFL	43.57 $\pm$ 0.55	40.88 $\pm$ 0.61	37.29 $\pm$ 0.51	34.98 $\pm$ 0.55
DPSIR-Fed	47.82 $\pm$ 0.44	45.23 $\pm$ 0.51	43.08 $\pm$ 0.39	40.47 $\pm$ 0.46
FedSS	50.93$\pm$0.37	48.24$\pm$0.43	45.72$\pm$0.35	43.08$\pm$0.41

7 Conclusion

In this paper, we proposed FedSS, a shareability-aware privacy-preserving federated learning framework for sports image recognition. FedSS distinguishes transferable sport semantics from client-specific visual context and introduces a SSSM to evaluate stage-wise model components according to sport discriminativeness, cross-client invariance, and privacy robustness. The resulting shareability scores guide private aggregation so that more transferable information is preserved while client-specific context remains local. Experiments on sport22 and sport100 demonstrate that FedSS consistently improves recognition accuracy and the privacy-utility trade-off under different heterogeneous and privacy settings. Further analyses show stronger cross-client semantic consistency, lower client-context predictability, and reduced clipping rate and distortion compared with representative DP-FL baselines. These results demonstrate the effectiveness of shareability-aware private federation for heterogeneous sports image recognition. Future work will extend FedSS to video and multimodal sports data, investigate naturally occurring cross-institutional data heterogeneity, and explore finer-grained adaptive shareability estimation.

Author contributions

Jiahui Lin conceived the study, designed the methodology, implemented the algorithm, conducted the experiments, analyzed the results, and drafted the manuscript. Jie Wu contributed to the experimental design, result interpretation, and manuscript revision. Ke Sun supervised the study, provided guidance on the research design, and revised the manuscript. Yuanyuan Huang contributed to manuscript review and revision. All authors read and approved the final manuscript.

Funding

This research received no external funding.

Data availability

The datasets analyzed in this study are publicly available in Kaggle repositories. The 100 Sports Image Classification dataset is available at <https://www.kaggle.com/datasets/gpiosenka/sports-classification>, and the 22 Sports Image Classification dataset is available at <https://www.kaggle.com/datasets/sheikhzaib/sports-image-image-classification>. The processed federated data partitions and experimental configuration files generated for this study are available from the corresponding authors upon reasonable request.

Ethics approval and consent to participate

Not applicable. This study used publicly available benchmark datasets and did not involve newly recruited human participants, human-subject experiments, or the collection of new personal data.

Consent for publication

Not applicable.

Competing interests

The authors declare no competing interests.

References

- [1] Shneiderman, B. Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human-Computer Interaction* **36**, 495–504 (2020). URL <https://doi.org/10.1080/10447318.2020.1741118>.
- [2] Garibay, O. O. *et al.* Six human-centered artificial intelligence grand challenges. *International Journal of Human-Computer Interaction* **39**, 391–437 (2023). URL <https://doi.org/10.1080/10447318.2022.2153320>.
- [3] Zhao, Z. *et al.* A survey of deep learning in sports applications: Perception, comprehension, and decision. *IEEE Transactions on Visualization and Computer Graphics* (2025).
- [4] Mateus, N. *et al.* Empowering the sports scientist with artificial intelligence in training, performance, and health management. *Sensors* **25**, 139 (2024).
- [5] Saltarella, M., Desolda, G., Lanzilotti, R. & Barletta, V. S. Translating privacy design principles into human-centered software lifecycle: A literature review. *International Journal of Human-Computer Interaction* **40**, 4465–4483 (2024). URL <https://doi.org/10.1080/10447318.2023.2219964>.

- [6] Liu, Q. & Wu, H. Privacy protection framework for human-machine interaction in advanced manufacturing empowered by artificial intelligence. *International Journal of Human-Computer Interaction* **0**, 1–16 (2025). URL <https://doi.org/10.1080/10447318.2025.2597548>.
- [7] McMahan, B., Moore, E., Ramage, D., Hampson, S. & y Arcas, B. A. *Communication-efficient learning of deep networks from decentralized data*. *Artificial intelligence and statistics*, 1273–1282 (PMLR, 2017).
- [8] Tian, L. *et al.* Federated optimization in heterogeneous networks. *arXiv: Learning, arXiv: Learning* (2018).
- [9] Fredrikson, M., Jha, S. & Ristenpart, T. *Model inversion attacks that exploit confidence information and basic countermeasures*. *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, 1322–1333 (2015).
- [10] Shokri, R., Stronati, M., Song, C. & Shmatikov, V. *Membership inference attacks against machine learning models*. *2017 IEEE symposium on security and privacy (SP)*, 3–18 (IEEE, 2017).
- [11] Zhu, L., Liu, Z. & Han, S. Deep leakage from gradients. *Advances in neural information processing systems* **32** (2019).
- [12] Geiping, J., Bauermeister, H., Dröge, H. & Moeller, M. Inverting gradients - how easy is it to break privacy in federated learning? *ArXiv abs/2003.14053* (2020). URL <https://api.semanticscholar.org/CorpusID:214728347>.
- [13] Noble, M., Bellet, A. & Dieuleveut, A. *Differentially private federated learning on heterogeneous data*. *International conference on artificial intelligence and statistics*, 10110–10145 (PMLR, 2022).
- [14] Geyer, R. C., Klein, T. & Nabi, M. Differentially private federated learning: A client level perspective. *arXiv preprint arXiv:1712.07557* (2017).
- [15] Wei, K. *et al.* Federated learning with differential privacy: Algorithms and performance analysis. *IEEE transactions on information forensics and security* **15**, 3454–3469 (2020).
- [16] Fu, J. *et al.* Differentially private federated learning: A systematic review. *arXiv preprint arXiv:2405.08299* (2024).
- [17] Cheng, A., Wang, P., Zhang, X. S. & Cheng, J. *Differentially private federated learning with local regularization and sparsification*. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10122–10131 (2022).

- [18] Shi, Y. *et al.* *Make landscape flatter in differentially private federated learning. Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 24552–24562 (2023).
- [19] Hu, R., Guo, Y. & Gong, Y. Federated learning with sparsified model perturbation: Improving accuracy under client-level differential privacy. *IEEE Transactions on Mobile Computing* **23**, 8242–8255 (2023).
- [20] Zhang, X., Xu, R., Yu, H., Zou, H. & Cui, P. *Gradient norm aware minimization seeks first-order flatness and improves generalization. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20247–20257 (2023).
- [21] Reis, F. J., Alaiti, R. K., Vallio, C. S. & Hespanhol, L. Artificial intelligence and machine learning approaches in sports: Concepts, applications, challenges, and future perspectives. *Brazilian journal of physical therapy* **28**, 101083 (2024).
- [22] Zhao, Y. *et al.* Federated learning with non-iid data. *Cornell University - arXiv, Cornell University - arXiv* (2018).
- [23] Li, Q., Diao, Y., Chen, Q. & He, B. *Federated learning on non-iid data silos: An experimental study. 2022 IEEE 38th International Conference on Data Engineering (ICDE)* (2022). URL <http://dx.doi.org/10.1109/icde53745.2022.00077>.
- [24] Zhao, Z. *et al.* Federated learning with non-iid data in wireless networks. *IEEE Transactions on Wireless Communications* 1927–1942 (2022). URL <http://dx.doi.org/10.1109/twc.2021.3108197>.
- [25] Li, X., Jiang, M., Zhang, X., Kamp, M. & Dou, Q. *Fedbn: Federated learning on non-iid features via local batch normalization. International Conference on Learning Representations* (2021).
- [26] Zhou, Z. *et al.* Model inversion attacks: A survey of approaches and countermeasures. *arXiv preprint arXiv:2411.10023* (2024).
- [27] Bai, L. *et al.* Membership inference attacks and defenses in federated learning: A survey. *ACM Computing Surveys* **57**, 1–35 (2024).
- [28] Zhao, J. C., Dabholkar, A., Sharma, A. & Bagchi, S. *Leak and learn: An attacker's cookbook to train using leaked data from federated learning. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12247–12256 (2024).
- [29] Wang, X., Wu, L. & Guan, Z. Graddiff: Gradient-based membership inference attacks against federated distillation with differential comparison. *Information Sciences* **658**, 120068 (2024).

- [30] Dwork, C., Roth, A. *et al.* The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science* **9**, 211–407 (2014).
- [31] El Ouadrhiri, A. & Abdelhadi, A. Differential privacy for deep and federated learning: A survey. *IEEE access* **10**, 22359–22380 (2022).
- [32] Demelius, L., Kern, R. & Trügler, A. Recent advances of differential privacy in centralized deep learning: A systematic survey. *ACM Computing Surveys* **57**, 1–28 (2025).
- [33] Chen, L., Ding, X., Bao, Z., Zhou, P. & Jin, H. Differentially private federated learning on non-iid data: Convergence analysis and adaptive optimization. *IEEE Transactions on Knowledge and Data Engineering* **36**, 4567–4581 (2024).
- [34] Andrew, G., Thakkar, O., McMahan, B. & Ramaswamy, S. Differentially private learning with adaptive clipping. *Advances in Neural Information Processing Systems* **34**, 17455–17466 (2021).
- [35] Xia, T. *et al.* Differentially private learning with per-sample adaptive clipping. *Proceedings of the AAAI Conference on Artificial Intelligence* **37**, 10444–10452 (2023).
- [36] Xue, R. *et al.* Differentially private federated learning with an adaptive noise mechanism. *IEEE Transactions on Information Forensics and Security* **19**, 74–87 (2023).
- [37] Li, Y., Yang, S., Ren, X., Shi, L. & Zhao, C. Multi-stage asynchronous federated learning with adaptive differential privacy. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **46**, 1243–1256 (2023).
- [38] Lai, J. *et al.* Sparsity-induced reparametrization for differentially private federated learning. *ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2311–2315 (IEEE, 2026).
- [39] Collins, L., Hassani, H., Mokhtari, A. & Shakkottai, S. *Exploiting shared representations for personalized federated learning. International conference on machine learning*, 2089–2099 (PMLR, 2021).
- [40] Chen, H.-Y. & Chao, W.-L. *On bridging generic and personalized federated learning for image classification. International Conference on Learning Representations* (2022).
- [41] Shen, Y., Zhou, Y. & Yu, L. *Cd 2-pfed: Cyclic distillation-guided channel decoupling for model personalization in federated learning. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10031–10040 (IEEE, 2022).

- [42] Kaggle. 22 sports image classification. <https://www.kaggle.com/datasets/sheikhzaib/sports-image-image-classification> (2022). Accessed: 2022.
- [43] Kaggle. 100 sports image classification. <https://www.kaggle.com/datasets/gpiosenska/sports-classification> (2023). Accessed: 2023-10-15.
- [44] McMahan, H. B., Ramage, D., Talwar, K. & Zhang, L. Learning differentially private recurrent language models. *arXiv preprint arXiv:1710.06963* (2017).
- [45] Sun, Y., Shen, L., Huang, T., Ding, L. & Tao, D. Fedspeed: Larger local interval, less communication round, and higher generalization accuracy. *arXiv preprint arXiv:2302.10429* (2023).