Domain-Aware Suppression and Aggregation for Federated DG ReID

Zhixi Yu[1*], Wei Liu[1*], Wenke Huang[2*], Bin Yang[2], Qian Bie[1], Guancheng Wan[3], Xin Xu[1†]

[1] School of Computer Science and Technology, Wuhan University of Science and Technology, Wuhan 430065 China
[2] National Engineering Research Center for Multimedia Software, School of Computer Science, Wuhan University, China
[3] University of California, Los Angeles
yuzhixi@wust.edu.cn, liuwei@wust.edu.cn, wenkehuang@whu.edu.cn,
yangbin_cv@whu.edu.cn, bieqian@wust.edu.cn, guanchengwan@whu.edu.cn ,xuxin@wust.edu.cn

Abstract

Federated domain generalization in person re-identification (FedDG-ReID) aims to learn a privacy-preserving server model from decentralized client source domains that generalizes to unseen domains. Existing approaches enhance the generalizability of the server model by increasing the diversity of client person data. However, these methods overlook that ReID model parameters are easily biased by client-specific data distributions, leading to the capture of excessive domain-specific identity information. Such identity information (*e.g.*, clothing style) struggles with identity information in unseen domains, thereby hindering the generalization ability of the server model. To address this, we propose a novel **Fed**DG-ReID framework, which mainly consists of Domain-aware Parameter **Sup**pression (DPS) and Domain-invariant **W**eighted **A**ggregation (DWA), called **FedSupWA**. Specifically, DPS adaptively attenuates the update magnitude of the parameters based on the fit of the parameters to the client's domain, encouraging the model to focus on more generalized domain-independent identity information, such as pedestrian contours, and other consistent information across domains. DWA enhances the server model's generalization by evaluating the effectiveness of the client model in maintaining the consistency of pedestrian identities to measure the importance of the learned domain-independent identity information and assigning greater aggregation weights to clients that contribute more generalized information. Extensive experiments demonstrate the effectiveness of FedSupWA, showing that it achieves state-of-the-art performance.

Introduction

Person re-identification (ReID) (Ye et al. 2019, 2021; Huang et al. 2019; Liu, Chang, and Shen 2020), which aims to retrieve the target person in cross-camera scenarios, stands as a crucial task in video intelligent surveillance systems (Li, Guo, and Wang 2021) and plays a critical role in smart-city design (*e.g.,* population mobility management) and public safety (*e.g.,* searching for a lost child or an escaped criminal). As research progresses, an increasing number of studies within the field of ReID focus on the problem of domain generalization (DG) (Zhou et al. 2022; Zhao et al. 2021; Dai et al. 2021; Xu et al. 2022; Liu et al. 2025). Research on domain generalization for person ReID focuses on constructing robust models that can effectively generalize to unseen domains by integrating training data from multiple source domains. Recent approaches (Zhou et al. 2020; Guo et al. 2020; Choi et al. 2021; Zhang et al. 2021) exploit multi-source training to enhance generalization by learning domain-invariant representations. While effective, these methods typically rely on centralized training over large-scale datasets, raising privacy concerns in sensitive, human-centric scenarios like ReID.

Therefore, this privacy concern has motivated researchers to explore federated learning (FL) (Ma et al. 2025; Huang et al. 2024, 2023a; Li, He, and Song 2021; Huang et al. 2023b; Huang, Ye, and Du 2022). FL is a distributed machine learning approach that aims to utilize data from multiple clients for collaborative model training while protecting data privacy. However, in the person ReID task, a sig-

*These authors contributed equally.
†Corresponding author.

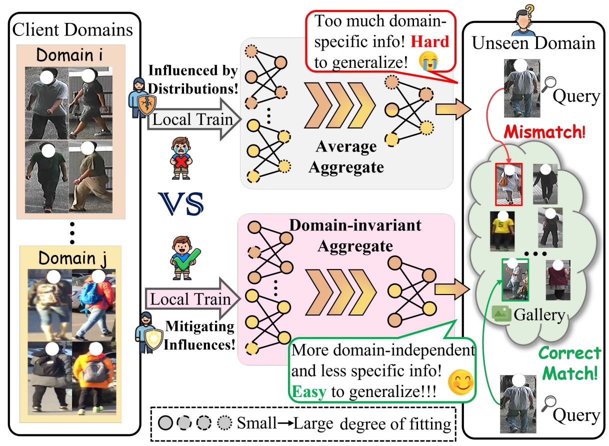

Figure 1: Problem illustration. Existing methods overlook that parameters are easily biased by client-specific data distributions, which leads to the learning of domain-specific identity information (*e.g.*, clothing style) and hinders generalization. In contrast, our approach suppresses this by promoting domain-invariant information learning (*e.g.*, pedestrian contours), improving performance on unseen domains.

nificant domain gap exists between client source domains and the unseen server domain. The gap arises from variations in person identities and differences in the shooting environments across client datasets. As a result, non-centrally trained models struggle to achieve good generalization performance on unseen domains. Existing FedDG-ReID approaches (Wu and Gong 2021; Yang et al. 2024) improve server generalization by leveraging the diversity of client person data. However, these methods tend to overlook the fact that the ReID model ***parameters are easily influenced by client-specific data distributions***, leading to the capture of excessive domain-specific identity information. Such domain-specific identity information (*e.g.*, clothing style) poorly applies to unseen domains, ultimately degrading generalization of the server model. As illustrated in Figure 1, excessive domain-specific identity information hinders the generalization of the server model.

To address the above problems, we propose a novel FedDG-ReID framework, called **FedSupWA**, which mainly consists of Domain-aware Parameter Suppression (DPS) and Domain-invariant Weighted Aggregation (DWA). Specifically, DPS mitigates the influence of client data distributions by dynamically attenuating the update magnitude of model parameters based on the degree of client model fitting. Since the server model aggregates knowledge from multiple clients, the difference between the client and server models reflects the extent to which the client model fits its local data distribution. When a parameter in the client model differs significantly from its counterpart in the server model, it suggests that the parameter deviates from the global trend and is more strongly fitted to the local data. Based on this insight, we quantify the degree of parameter fitting by measuring the difference between the client's local model and the server's global model. Parameter updates are then attenuated adaptively according to this degree of fit: parameters with larger differences—indicating stronger fitting to local data—are more heavily suppressed. This difference-driven adaptive attenuation mechanism enables DPS to effectively reduce the learning of domain-specific identity information and mitigate the impact of client data distributions on the aggregated model.

In conjunction with DPS, DWA dynamically adjusts aggregation weights based on the importance of domain-independent identity information captured by client models, thereby enhancing the generalization of the server model. Since DPS suppresses the learning of domain-specific identity information, improvements in identity consistency primarily reflect the acquisition of domain-invariant information, such as pedestrian contours, and other consistent information across domains. DWA quantifies this importance by monitoring changes in client identity consistency, assigning higher weights to clients that contribute more generalizable information. This dynamic aggregation enables the server model to integrate domain-independent identity information more effectively, thereby improving its performance on unseen domains. Comprehensive experiments validate the effectiveness of our framework, showing robust generalization across person ReID datasets. Our contributions are summarized as follows:

- **Framework Contribution.** We propose a novel FedDG-ReID framework, **FedSupWA**, which mitigates the influence of client-specific data distributions by attenuating their update magnitudes and dynamically aggregating client models based on their contributions to domain-invariant identity information learning, thereby enhancing generalization to unseen domains.
- **Technical Contribution.** We introduce DPS and DWA: DPS adaptively attenuates parameter updates based on their degree of fitting, suppressing domain-specific identity information while promoting domain-invariant identity information learning. Building on this, DWA dynamically aggregates client models according to the importance of domain-invariant identity information they capture, further mitigating the effect of client-specific data distributions on the server model.
- **Empirical Contribution.** Our approach consistently achieves state-of-the-art performance across all three standard benchmarks under unseen domain settings, significantly outperforming existing FedDG-ReID methods.

Related Work

Domain Generalization for Re-ID

Domain generalization aims to learn models from multiple seen source domains that can generalize to unseen domains. Recently, researchers have proposed several methods to solve the domain generalization problem in person reidentification (ReID). (Song et al. 2019) introduces the problem of domain generalization in person ReID and proposes the Domain-Invariant Mapping Network, which employs a memory bank to learn domain-invariant features. (Jin et al. 2020) applies instance normalization to separate identity-relevant and identity-irrelevant features. The model's generalization ability is strengthened by recovering discriminative information from the identity-relevant features. In contrast, a full-scale network (Zhou et al. 2020) is designed to learn discriminative features directly. On the other hand, (Zhuang et al. 2020b) utilizes a cross-camera invariant representation using camera-based batch normalization, which mitigates feature bias during cross-camera transfer and improves the model's scene adaptability. In parallel, several studies have explored meta-learning approaches for DG person ReID. (Zhao et al. 2021) enhances model generalization by integrating a long-term memory mechanism with the rapid adaptation capabilities of meta-learning, leveraging interpolated features. Similarly, (Choi et al. 2021) employs meta-learning to simulate failure cases in Batch Normalization (BN) and Instance Normalization (IN) and dynamically adjusts the reliance on normalization layers accordingly.

Despite the success of these approaches, most of them require aggregating data from different source domains, which certainly increases the risk of privacy leakage, especially for person ReID datasets containing a large amount of personal information. Therefore, this paper attempts to address the task of person ReID domain generalization in a federated learning scenario, which improves the generalization ability of the model while preserving privacy.

Federated Learning

Federated learning (McMahan et al. 2017; Karimireddy et al. 2020; Liu et al. 2021) is a solution for decentralized training data (Kou et al. 2025b,c,a) in machine learning, and its key mechanism lies in model training locally on the client and global aggregation of decentralized model parameters through the server. FedAvg (McMahan et al. 2017) is one of the earliest federated optimization methods. It performs global updates by weighted averaging of model parameters trained locally on client devices. While effective under consistent data distributions, it suffers from unstable aggregation and performance degradation in the presence of data heterogeneity (non-IID conditions). To address data heterogeneity in federated learning, FedProx (Li et al. 2018) introduces a proximal regularization term that constrains the magnitude of local updates. It helps prevent local models from deviating excessively from the global model. MOON (Li, He, and Song 2021) mitigates personalization bias by aligning local model representations with the global model through model-level contrastive learning, thereby reducing performance degradation. SCAFFOLD (Karimireddy et al. 2020) tackles client drift caused by non-IID data by introducing control variates to adjust the direction of local updates. While these approaches improve robustness in non-IID settings, they treat the client model as a whole and overlook the fact that client-specific data distributions affect different parameters to varying degrees.

Compared to the closed-set classification problem, federated person ReID faces more severe open-set challenges: non-overlapping identity classes in the training and test sets, and significant cross-domain distributional differences in data between clients. FedPav (Zhuang et al. 2020a) proposes federated person ReID and designs a global aggregation strategy based on the cosine distance of classifier logits. (Wu and Gong 2021) first explicitly introduces the definition of FedDG-ReID by considering each known source domain as a client and obtaining models with the capability of generalizing to unseen domains by aggregating client models. In order to improve the generalization ability of the model, (Yang et al. 2024) and (Liu, Ye, and Du 2024) improve the model's adaptability to unseen data by increasing the diversity of client data domains and creating local virtual domains, respectively. However, existing methods often neglect that model parameters are easily biased by client-specific data distributions, which results in the encoding of excessive domain-specific identity information that generalizes poorly to unseen domains, ultimately degrading server performance. To address this, our approach enhances generalization by attenuating parameter updates based on their degree of fit and dynamically aggregating client models according to the importance of the domain-invariant identity information they capture.

Methodology

Problem Definition

Assume that there are N private datasets (*e.g.*, MSMT17 (Wei et al. 2018), CUHK03 (Li et al. 2014)) with labels to serve as the source domains for the clients, denoted as $\mathcal{S} = \{\mathcal{D}_1, \mathcal{D}_2, \ldots, \mathcal{D}_N\}$. Each dataset is considered an individual client. In this paper, "domain" and "client" are interchangeable. Each domain $\mathcal{D}_i = \{X_i, Y_i\}$ consists of M_i training samples, where X_i represents the images and Y_i the corresponding identity labels. To protect the privacy of client personal data, the FedDG-ReID models are trained through a co-training mechanism without sharing local samples of clients. The training procedure for federated person ReID during a communication round can be briefly summarized in the following four steps:

- **Distribution**: The server dispatches the aggregated global backbone to each client for subsequent local training. In particular, the server will distribute a pre-initialized backbone in the first round of communication.
- **Local training**: Upon receipt of the aggregated global backbone from the server, each client replaces its local backbone with the global counterpart and proceeds to train the model on its proprietary dataset.
- **Client-to-server updating**: Following local training completion, each client transmits the trained backbone network to the server while retaining the local classifier.
- **Model aggregation**: The server aggregates the clients' backbone networks based on a predefined policy and updates the global backbone.

Following numerous rounds of iterative communication, the global model achieves convergence. In the subsequent sections, we delineate in detail our proposed methodologies: Domain-aware Parameter Suppression and Domain-invariant Weighted Aggregation.

Domain-aware Parameter Suppression

To mitigate the impact of client data-specific data distributions on model parameters, we propose Domain-aware Parameter Suppression (DPS). In this strategy, we first quantify the degree of parameter fitting during local training by measuring the discrepancy between the client's local parameters and the server-aggregated model parameters. To enable update attenuation, the computed differences are normalized into fit coefficients, where larger discrepancies indicate higher degrees of fitting and correspond to lower fit coefficients. Ultimately, the update magnitude of each parameter is attenuated according to its fitting coefficient, effectively reducing the impact of client-specific data distributions on model parameters. To elucidate the methodology for assessing the model's fit and the subsequent adaptive updates, we exemplify the parameters of the n-th layer of the server during the k-th round of communication. In the ensuing discussion, we denote the parameters corresponding to the n-th layer of the model as "model parameters" for brevity.

Quantify the Fitting Degree. We begin by computing the parameter difference $\Delta W_{i,k}$ between the server model Wk^* and the client model $W_{i,k}$ (absolute difference), defined as:

$$\Delta W_{i,k} = |W_k^* - W_{i,k}| \,, \tag{1}$$

where W_k^* denotes the server parameters during the k-th round of communication, $W_{i,k}$ signifies the model parameters after the (k-1)-th round of communication of the i-th

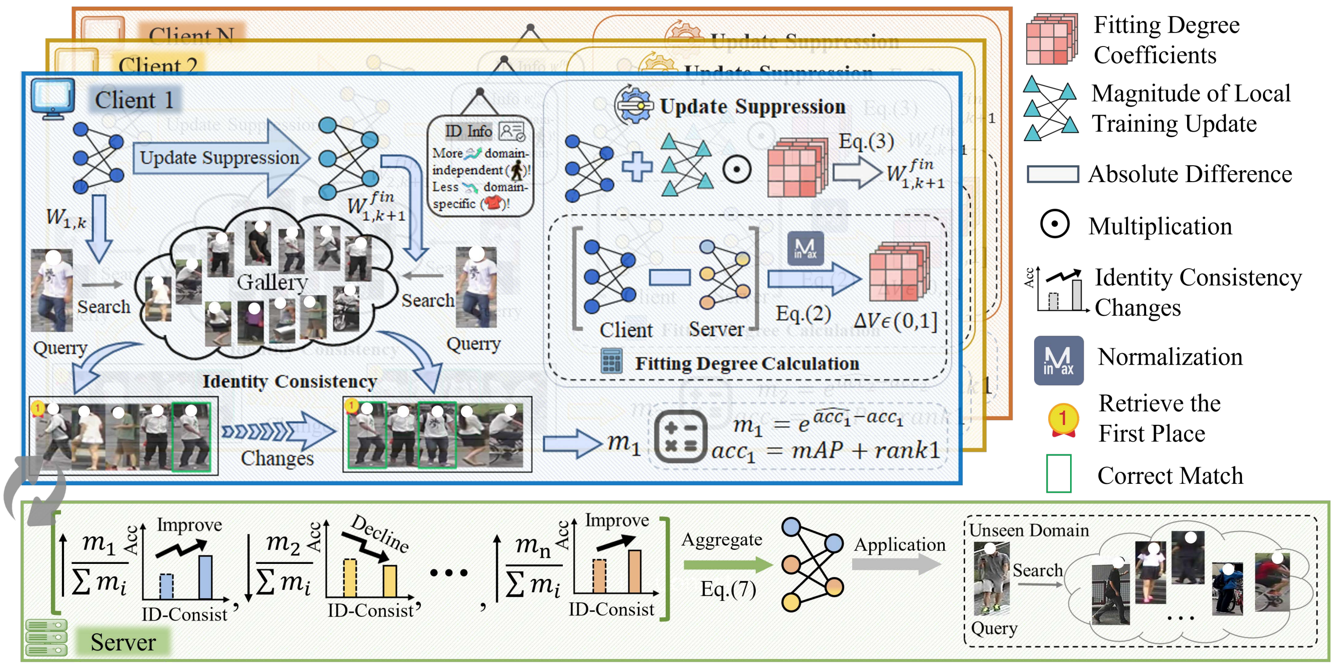

Figure 2: Architecture illustration of FedSupWA. The figure illustrates the process of parameter update suppression and identity-consistency-guided aggregation, using Client 1 as an example. For clarity, the corresponding processes of other clients are overlaid in the background using different colors. While mAP and $rank1$ are evaluated on the client during training, redundant computation on the server is avoided during inference.

client. Here, $\Delta W_{i,k} = \{\Delta w_1, \Delta w_2, \ldots, \Delta w_j, \ldots, \Delta w_m\}$, Δw_j encapsulates the j-th parameter difference, and m denotes the number of parameters.

Given the computed ΔW_k, we evaluate the degree to which the client model fits its local data. Specifically, the server model, obtained by aggregating multiple clients, reflects the global data distribution. A large discrepancy between a client parameter and the global model indicates strong deviation from the global trend and a high degree of fitting to the local data distributions. Conversely, a smaller difference implies better alignment with the global distribution and a lower fit. Since parameters vary in scale, we normalize Δw_j as follows to obtain a consistent measure of their fitting degrees:

$$\Delta v_j = 1 - \left(\frac{\Delta w_j - min(\Delta W_k)}{max(\Delta W_k) - min(\Delta W_k) + \epsilon}\right), \quad (2)$$

where ϵ is set to 0.0001 to prevent instability of division by zero. The terms min and max denote the minimum and maximum values, respectively, while Δv_j represents the normalized parameter fitting coefficient.

According to Eq. 2, an increase in the parameter difference Δw_j results in a decrease of the fitting coefficient Δv_j. This quantitative relationship implies that larger discrepancies between client and server parameters correspond to stronger local fitting and thus lower fitting coefficients. Conversely, smaller differences indicate better alignment with the global model and higher fitting coefficients.

Update Suppression. Based on the fitting coefficients, we attenuate client parameter updates to varying degrees. Specifically, in the k-th communication round, the server model is first synchronized to the client after computing the fitting coefficients. The client then performs local training to obtain its updated model. By scaling the pre-update parameter changes with their respective fitting coefficients, we derive the adjusted update magnitudes. The final client model is then obtained by adding these altered updates to the client's pre-update parameters. The detailed formulation of the updated client model is presented as follows:

$$W_{i,k+1}^{fin} = W_{i,k} + \Delta V \odot (W_{i,k+1} - W_{i,k}), \quad (3)$$

$$W_{i,k} = W_k^*, \quad (4)$$

where $W_{i,k+1}$ signifies the client model after training on local data, $W_{i,k+1}^{fin}$ refers to the final model of client i after the (k+1)-th round, $W_{i,k}$ denotes the client model after overwriting it with the server model, and $\odot$ represents the Hadamard product. Eq. 4 indicates that the server distributes the model to the client. From Eq. 2, it is clear that $\Delta V \in (0, 1]$, the update magnitude of the client model is attenuated to varying degrees according to its fitting coefficient.

Domain-aware Parameter Suppression strategy suppresses the learning of domain-specific identity information while promoting the acquisition of domain-independent identity information by dynamically reducing the magnitude of parameter updates, thereby laying the foundation for subsequent server aggregation.

Algorithm 1: FedSupWA

Require: N: client number, E_g: global epochs, E_l: local epochs, η: learning rate, W_0^*: initial parameters
Ensure: $W_{E_g}^*$: weights of global model

function $\textsc{Server}(W_0^*)$
  for round $k = 0$ **to** $E_g - 1$ **do**
    for each client $i \in N$ **concurrently do**
      $W_{i,k+1}, m_{i,k+1} \leftarrow \textsc{ClientSide}(W_k^*, i, k+1)$
    end for
    $W_{k+1}^* \leftarrow \sum_{i \in N} \frac{m_{i,k+1}}{\sum_{j \in N} m_{j,k+1}} W_{i,k+1}$
  end for
  return $W_{E_g}^*$
end function
function $\textsc{ClientSide}(W_k^*, i, k+1)$
  if $k > 0$ **then**
    $\Delta W_{k+1} = |W_k^* - W_{i,k}|$
    $\Delta V_{k+1} \leftarrow$(transform ΔW_{k+1} via Eq. 2)
  end if
  $W_{i,k} \leftarrow W_k^*$
  for each local epoch $e = 0 \rightarrow E_l - 1$ **do**
    $(W_{i,k+1}, V_{i,k+1}) \leftarrow (W_{i,k}, V_{i,k}) - \eta \nabla \mathcal{L}\big((W_{i,k}, V_{i,k})\big)$
  end for
  $W_{i,k+1} \leftarrow W_{i,k} + \Delta V_{k+1} \cdot (W_{i,k+1} - W_{i,k})$
  $m_{i,k+1} \leftarrow$(calculate Δacc via Eq. 5)
  return $W_{i,k+1}, m_{i,k+1}$
end function

Domain-invariant Weighted Aggregation

While DPS curbs the learning of domain-specific identity information and enhances domain-invariant identity information, uniformly weighting models that capture varying domain-invariant identity information can still constrain the server model's generalization capacity. Therefore, we propose the Domain-invariant Weighted Aggregation (DWA). Since DPS suppresses domain-specific feature learning, changes in the client model's pedestrian identity consistency primarily reflect the importance of its domain-invariant identity information. Building on this, DWA quantifies the significance of domain-invariant identity information by measuring changes in the client model's identity consistency. Specifically, the strategy employs the commonly used mAP and Rank-1 metrics in ReID as measures of identity consistency. Initially, the mAP and Rank-1 metrics of the model on local data before and after update suppression will be calculated for each client. An increase in mAP or Rank-1 signifies improved learning of domain-invariant identity information, whereas a decrease indicates the opposite. Moreover, larger performance gains correspond to greater significance of the learned domain-independent identity information. Clients exhibiting greater improvements in domain-invariant identity learning are assigned higher aggregation weights, while those with smaller gains receive lower weights. The aggregation weights are computed as follows:

$$m_i = e^{\widehat{acc}_i - acc_i}, \tag{5}$$

$$acc_i = mAP_i + rank_i^1, \tag{6}$$

where $\widehat{acc}_i$ and acc_i denote the accuracy before and after the differential update of the i-th client model, including mAP and Rank-1, respectively, as in Eq. 6. The calculation of $\widehat{acc}_i$ is the same as acc_i. In Eq. 5, the natural logarithm is applied to ensure numerical stability by preventing the denominator from approaching zero during aggregation.

According to the variation of model performance of each client model, we can calculate the weights at the time of aggregation by Eq. 6, and the server aggregates according to the weights to get the global model:

$$W_{k+1}^* = \sum_{i=1}^{N} \frac{m_i}{\sum m_i} W_{i,k+1}^{fin}, \tag{7}$$

where W_{k+1}^* denotes the server model after aggregation, $W_{i,k+1}^{fin}$ denotes the client model, m_i represents the aggregation weights, and N denotes the number of clients.

DWA dynamically aggregates client models based on changes in performance metrics before and after local updates. In this manner, the variation in domain-invariant identity information guides the aggregation weights, reducing the influence of clients with limited generalizable information and enhancing the global model's overall performance.

Experiments

Experiment Setups

Datasets. We conduct experiments on four large-scale ReID benchmarks: Market-1501 (M) (Zheng et al. 2015), CUHK02 (C2) (Li and Wang 2013), MSMT17 (MS) (Wei et al. 2018), and CUHK03 (C3) (Li et al. 2014). Following the experimental setup of (Yang et al. 2024), three datasets are used as training sources, each assigned to a separate local client (one full dataset per client), while the remaining dataset serves as an independent unseen test domain.

Implementation Details. We adopt ResNet50 and ViT (pretrained on ImageNet) as backbone networks. Due to computational constraints, the batch size is set to 64 for ResNet50 and 16 for ViT-based models. Training is performed over 60 communication rounds, with each client conducting 200 local iterations per round. All models are implemented in Python 3.9 using PyTorch 2.1.0 and trained on a single RTX 3090 GPU. Performance is evaluated using mean Average Precision (mAP) and Rank-1 accuracy.

Comparison with State-of-the-arts Methods

We compare our methods with the state-of-the-art in Table 1. The methods are grouped into four categories: (1) Federated Learning methods, including SCAFFOLD (Karimireddy et al. 2020), MOON (Li, He, and Song 2021), and FedProx (Li et al. 2020), which are trained on ReID datasets to evaluate the applicability of standard federated learning to the FedDG-ReID task; (2) Domain Generalization methods, such as SNR (Jin et al. 2020), MixStyle (Zhou et al. 2021), and CrossStyle (Tang et al. 2021), where each client independently applies the domain generalization strategy during local training. (3) Federated ReID methods, including

Category	Methods	Venue	Market1501		MSMT17		CUHK03		Average	
			mAP	Rank-1	mAP	Rank-1	mAP	Rank-1	mAP	Rank-1
ResNet50 Backbone										
Federated Learning	SCAFFOLD	ICML'20	26.0	50.5	5.3	15.8	22.9	26.0	18.1	30.8
	MOON	CVPR'21	26.8	51.1	4.8	14.5	20.9	22.5	17.5	29.4
	FedProx	MLSys'20	29.3	53.8	5.8	17.4	19.1	17.7	18.1	29.6
Domain Generalization	SNR	CVPR'20	32.7	59.4	5.1	15.3	28.5	30.0	22.1	34.9
	MixStyle	ICLR'21	31.2	53.5	5.5	16.0	28.6	31.5	21.8	33.7
	CrossStyle	ICCV'21	35.5	59.6	4.6	14.0	27.8	28.0	22.6	33.9
Fed ReID	FedPav	MM'20	25.4	49.4	5.2	15.5	22.5	24.3	17.7	29.7
	FedReID	AAAI'21	30.1	53.7	4.5	13.7	26.4	26.5	20.3	31.3
FedDG-ReID	DACS	AAAI'24	36.3	61.2	10.4	27.5	30.7	34.1	25.8	40.9
	SSCU	MM'25	39.5	66.4	11.9	32.3	32.8	34.1	28.1	44.3
	FedSupWA	—	**42.0**	**68.9**	**12.4**	**33.4**	**36.1**	**36.6**	**30.2**	**46.3**
ViT Backbone										
Fed ReID	FedPav	MM'20	32.5	56.0	13.6	32.6	25.1	24.0	23.7	37.5
Domain Generalization	CrossStyle	ICCV'21	30.9	55.1	16.8	39.9	19.2	18.2	22.3	32.0
FedDG-ReID	DACS	AAAI'24	38.3	63.7	17.3	39.9	29.1	28.8	28.2	44.1
	SSCU	MM'25	46.3	69.6	20.0	**43.7**	31.5	31.6	32.6	48.3
	FedSupWA	—	**48.0**	**71.4**	**20.1**	**43.7**	**34.9**	**35.0**	**34.3**	**50.0**

Table 1: Comparison of unseen domain performance with state-of-the-art methods. We conduct experiments by rotating among Market1501, MSMT17, and CUHK03, using one dataset as the test set while training on the remaining three in each round.

Attributes		Market1501		MSMT17		CUHK03	
DPS	DWA	mAP	R1	mAP	R1	mAP	R1
×	×	39.5	66.4	11.9	32.3	32.8	34.1
✓	×	40.8	68.3	12.2	32.6	30.6	31.1
×	✓	40.7	67.7	12.3	32.9	35.1	36.1
✓	✓	**42.0**	**68.9**	**12.4**	**33.4**	**36.1**	**36.6**

Table 2: Ablation study of key components. DPS: Domain-aware Parameter Suppression, DWA: Domain-invariant Weighted Aggregation, R1: Rank-1.

FedPav (Zhuang et al. 2020a) and FedReID (Wu and Gong 2021); (4) Federated Domain Generalization ReID, including DACS (Yang et al. 2024) and SSCU (Xu et al. 2025). Across all evaluation settings, our method consistently outperforms existing methods. As shown in Table 1, it achieves an mAP of 36.1% (↑ **+3.4%** vs. 32.8% for SSCU) and a Rank-1 accuracy of 36.6% (↑ **+2.5%** vs. 34.1% for both SSCU and DACS) on the CUHK03. Notably, the generalizability of our method improves significantly when ViT is adopted as the backbone. On the Market1501 test set, it achieves 44.2% mAP, substantially outperforming FedPav (32.5%), CrossStyle (30.9%), and the recent DACS (38.3%), marking a substantial performance gain.

Ablation Study

To validate the effectiveness of our proposed DPS and DWA modules, we conduct multiple ablation studies to isolate their individual contributions. Given the critical role of aggregation in federated learning, we replace DWA with domain-size-weighted averaging to assess the impact of our dynamic aggregation strategy. Table 2 compares the performance of Baseline and Baseline+DPS, demonstrating that DPS effectively enhances generalization to unseen domains by dynamically adjusting update magnitudes based on parameter fitting degrees. The comparison between Baseline and Baseline+DWA further shows that dynamic aggregation guided by changes in identity consistency boosts server model performance. Importantly, combining both strategies yields the greatest improvement over the baseline, confirming their complementary benefits.

Further Study

Comparison of Different Aggregation Strategies. We compare different aggregation strategies, where "mAP+R1", "mAP", and "R1" denote metric-guided aggregation based on the corresponding evaluation indicators, and "Avg" refers to average aggregation weighted by domain size. As shown in Table 3, all metric-guided strategies (mAP, R1, and mAP+R1) outperform the "Avg" baseline, demonstrating the effectiveness of performance-based weighting. Notably, the "mAP+R1" strategy achieves the best performance, underscoring the advantage of leveraging multiple metrics over single-indicator guidance.

Comparison on More Client Domains. In Table 4, we evaluate the generalization ability of our method with an increased number of source domains. Specifically, we use 9 datasets: MSMT17, DukeMTMC (Zheng, Zheng, and Yang

Methods	Market1501		MSMT17		CUHK03	
	mAP	R1	mAP	R1	mAP	R1
Avg	40.8	68.3	12.2	32.6	30.6	31.1
mAP	41.6	68.3	12.2	32.7	35.9	36.3
R1	41.8	68.4	12.3	33.1	35.3	36.4
mAP+R1	**42.0**	**68.9**	**12.4**	**33.4**	**36.1**	**36.6**

Table 3: Comparison of different aggregation strategies. R1: Rank-1, Avg: aggregation based on domain size.

Methods	Market1501		MSMT17		CUHK03	
	mAP	R1	mAP	R1	mAP	R1
MOON	27.8	53.6	6.6	19.1	13.4	12.4
FedReID	33.8	59.0	7.8	22.0	20.7	19.6
FedPav	30.5	55.1	8.1	22.9	18.1	17.9
DACS	36.6	64.5	13.4	35.3	22.5	22.3
SSCU	40.5	68.5	14.0	36.2	24.4	24.8
FedSupWA	**40.9**	**69.0**	**14.3**	**37.0**	**25.0**	**25.1**

Table 4: Comparison of the performance in more client domain settings. R1 indicates Rank-1.

2017), Market1501, CUHK03, PRID (Hirzer et al. 2011), CUHK01 (Li, Zhao, and Wang 2012), VIPeR (Gray and Tao 2008), 3DPes (Baltieri, Vezzani, and Cucchiara 2011), and iLIDS (Wang et al. 2014), and designate Market-1501, MSMT17, and CUHK03 as unseen target domains. Even under the multi-source setting, our method consistently outperforms existing state-of-the-art approaches.

Performance Evolution Comparison

We compare the evolution of mAP and Rank-1 accuracy across 60 communication rounds on CUHK03 for FedPav, SSCU, FedReID, FedProx, DACS, CrossStyle, MOON, and our proposed method. As shown in Figure 3, our approach consistently outperforms all baselines from early training. Notably, FedSupWA (red) exhibits a remarkable reversal trajectory: starting from a low baseline of 5%in both metrics at epoch 0, it surpasses all competing methods by epoch 10 (reaching 30% mAP), breaks through the 30% threshold by epoch 25 (achieving 34% mAP and 35% Rank-1), and ultimately reaches 36% mAP and 36% Rank-1 at epoch 60. These substantial gains highlight the effectiveness of our DPS and DWA in improving generalization performance.

Visualization

We visualize the T-SNE embeddings at different training stages in Figure 4. For the "MS+C2+M" setting across epochs 10 to 60, the feature distributions of C2 (purple), M (orange), and MS (green) show a progressive convergence into more cohesive clusters as training advances. Notably, C2 and M exhibit accelerated alignment after epoch 30, and all datasets show reduced dispersion in later epochs. As highlighted in the magnified local views, similar samples increasingly form tighter clusters over time.

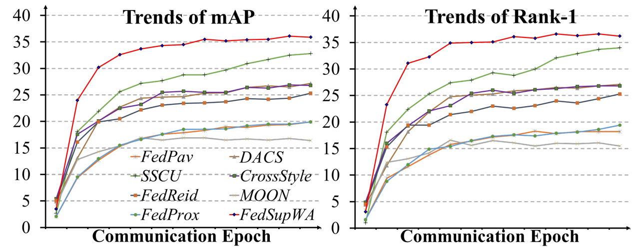

Figure 3: Comparison of convergence in mAP and Rank-1 in the selected methods. Please zoom in for the best view.

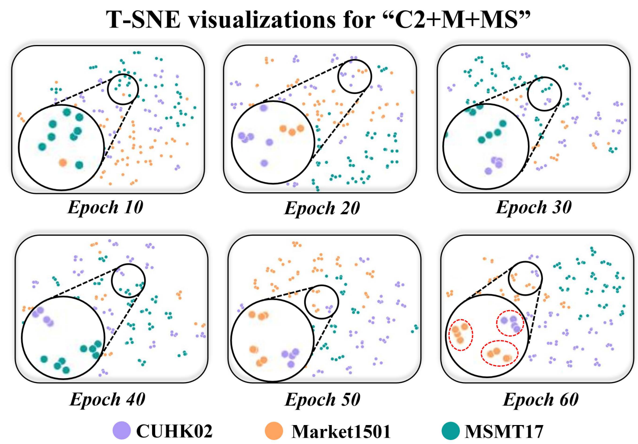

Figure 4: T-SNE visualizations of feature distribution on CUHK02 (C2), Market1501 (M), and MSMT17 (MS).

Conclusion

In this paper, we propose a novel federated generalized person ReID method, **FedSupWA**, which consists of Domain-aware Parameter Suppression (DPS) and Domain-invariant Weighted Aggregation (DWA). Specifically, DPS initially measures the fitting degree of parameters by calculating the difference between the client and server model, which then attenuates parameter updates based on the fit. Attenuating the updates of parameters, DPS inhibits the model's ability to capture domain-specific identity information and promotes the acquisition of shared knowledge. Since DPS inhibits the model's acquisition of domain-independent identity information, DWA computes the effectiveness of the client model in maintaining the consistency of pedestrian identities before and after the model's adaptive update to represent the change in the model's domain-independent identity information and subsequently aggregates the client models. By integrating these strategies, FedSupWA effectively mitigates the influence of client-specific data distributions on local models, thereby enhancing the generalization ability of the aggregated server model to unseen domains. We conducted extensive comparative experiments to demonstrate that our method outperforms existing state-of-the-art methods across different backbones (ResNet50 and ViT), and performed ablation studies to validate the effectiveness of the proposed DPS and DWA modules.

Acknowledgments

This work was supported by the Natural Science Foundation of China (62376201, 62501428, and 623B2080), Hubei Provincial Science & Technology Talent Enterprise Services Program (2025DJB059), and Hubei Provincial Special Fund for Central-Guided Local S&T Development (2025CSA017).

References

Baltieri, D.; Vezzani, R.; and Cucchiara, R. 2011. 3dpes: 3d people dataset for surveillance and forensics. In *Proceedings of the 2011 joint ACM workshop on Human gesture and behavior understanding*, 59–64.

Choi, S.; Kim, T.; Jeong, M.; Park, H.; and Kim, C. 2021. Meta batch-instance normalization for generalizable person re-identification. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 3425–3435.

Dai, Y.; Li, X.; Liu, J.; Tong, Z.; and Duan, L.-Y. 2021. Generalizable person re-identification with relevance-aware mixture of experts. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16145–16154.

Gray, D.; and Tao, H. 2008. Viewpoint invariant pedestrian recognition with an ensemble of localized features. In *European conference on computer vision*, 262–275. Springer.

Guo, J.; Zhu, X.; Zhao, C.; Cao, D.; Lei, Z.; and Li, S. Z. 2020. Learning meta face recognition in unseen domains. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6163–6172.

Hirzer, M.; Beleznai, C.; Roth, P. M.; and Bischof, H. 2011. Person re-identification by descriptive and discriminative classification. In *Scandinavian conference on Image analysis*, 91–102. Springer.

Huang, G.; Liu, Z.; Pleiss, G.; Van Der Maaten, L.; and Weinberger, K. Q. 2019. Convolutional networks with dense connectivity. *IEEE transactions on pattern analysis and machine intelligence*, 44(12): 8704–8716.

Huang, W.; Ye, M.; and Du, B. 2022. Learn from others and be yourself in heterogeneous federated learning. In *CVPR*.

Huang, W.; Ye, M.; Shi, Z.; and Du, B. 2023a. Generalizable Heterogeneous Federated Cross-Correlation and Instance Similarity Learning. *TPAMI*.

Huang, W.; Ye, M.; Shi, Z.; Li, H.; and Du, B. 2023b. Rethinking Federated Learning with Domain Shift: A Prototype View. In *CVPR*.

Huang, W.; Ye, M.; Shi, Z.; Wan, G.; Li, H.; Du, B.; and Yang, Q. 2024. A Federated Learning for Generalization, Robustness, Fairness: A Survey and Benchmark. *TPAMI*.

Jin, X.; Lan, C.; Zeng, W.; Chen, Z.; and Zhang, L. 2020. Style normalization and restitution for generalizable person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3143–3152.

Karimireddy, S. P.; Kale, S.; Mohri, M.; Reddi, S.; Stich, S.; and Suresh, A. T. 2020. Scaffold: Stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*, 5132–5143.

Kou, Z.; Qin, S.; Wang, H.; Xie, M.; Chen, S.; Jia, Y.; Liu, T.; Sugiyama, M.; and Geng, X. 2025a. Label Distribution Learning with Biased Annotations by Learning Multi-Label Representation. Main Track.

Kou, Z.; Wang, J.; Jia, Y.; Liu, B.; and Geng, X. 2025b. Instance-Dependent Inaccurate Label Distribution Learning. *IEEE Transactions on Neural Networks and Learning Systems*, 36(1): 1425–1437.

Kou, Z.; Xuan, H.; Zhu, J.; Wang, H.; Xie, M.-k.; Wang, C.; Wang, J.; Jia, Y.; and Geng, X. 2025c. Tail-Aware Reconstruction of Incomplete Label Distributions with Low-Rank and Sparse Modeling. *IEEE Transactions on Circuits and Systems for Video Technology*, 1–1.

Li, K.; Guo, D.; and Wang, M. 2021. Proposal-free video grounding with contextual pyramid network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 1902–1910.

Li, Q.; He, B.; and Song, D. 2021. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10713–10722.

Li, T.; Sahu, A. K.; Sanjabi, M.; Zaheer, M.; Talwalkar, A.; and Smith, V. 2018. On the convergence of federated optimization in heterogeneous networks. *arXiv preprint arXiv:1812.06127*, 6.

Li, T.; Sahu, A. K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; and Smith, V. 2020. Federated optimization in heterogeneous networks. *Proceedings of Machine learning and systems*, 2: 429–450.

Li, W.; and Wang, X. 2013. Locally aligned feature transforms across views. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3594–3601.

Li, W.; Zhao, R.; and Wang, X. 2012. Human reidentification with transferred metric learning. In *Asian conference on computer vision*, 31–44. Springer.

Li, W.; Zhao, R.; Xiao, T.; and Wang, X. 2014. Deepreid: Deep filter pairing neural network for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 152–159.

Liu, C.; Chang, X.; and Shen, Y.-D. 2020. Unity style transfer for person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6887–6896.

Liu, F.; Ye, M.; and Du, B. 2024. Domain generalized federated learning for Person Re-identification. *Computer Vision and Image Understanding*, 241: 103969.

Liu, Q.; Chen, C.; Qin, J.; Dou, Q.; and Heng, P.-A. 2021. Feddg: Federated domain generalization on medical image segmentation via episodic learning in continuous frequency space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1013–1023.

Liu, W.; Xu, X.; Chang, H.; Yuan, X.; and Wang, Z. 2025. Mix-modality person re-identification: A new and practical paradigm. *ACM Transactions on Multimedia Computing, Communications and Applications*, 21(4): 1–21.

Ma, Y.; Dai, W.; Huang, W.; and Chen, J. 2025. Geometric Knowledge-Guided Localized Global Distribution Alignment for Federated Learning. In *CVPR*.

McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, 1273–1282.

Song, J.; Yang, Y.; Song, Y.-Z.; Xiang, T.; and Hospedales, T. M. 2019. Generalizable person re-identification by domain-invariant mapping network. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 719–728.

Tang, Z.; Gao, Y.; Zhu, Y.; Zhang, Z.; Li, M.; and Metaxas, D. N. 2021. Crossnorm and selfnorm for generalization under distribution shifts. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 52–61.

Wang, T.; Gong, S.; Zhu, X.; and Wang, S. 2014. Person re-identification by video ranking. In *European conference on computer vision*, 688–703. Springer.

Wei, L.; Zhang, S.; Gao, W.; and Tian, Q. 2018. Person transfer gan to bridge domain gap for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 79–88.

Wu, G.; and Gong, S. 2021. Decentralised learning from independent multi-domain labels for person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 2898–2906.

Xu, X.; Liu, W.; Wang, Z.; Hu, R.; and Tian, Q. 2022. Towards generalizable person re-identification with a bi-stream generative model. *Pattern Recognition*, 132: 108954.

Xu, X.; Ren, C.; Liu, W.; Huang, W.; Yang, B.; Yu, Z.; and Jiang, K. 2025. Positive Style Accumulation: A Style Screening and Continuous Utilization Framework for Federated DG-ReID. *arXiv preprint arXiv:2507.16238*.

Yang, F.; Zhong, Z.; Luo, Z.; He, Y.; Li, S.; and Sebe, N. 2024. Diversity-authenticity co-constrained stylization for federated domain generalization in person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 6477–6485.

Ye, M.; Cheng, Y.; Lan, X.; and Zhu, H. 2019. Improving night-time pedestrian retrieval with distribution alignment and contextual distance. *IEEE Transactions on Industrial Informatics*, 16(1): 615–624.

Ye, M.; Shen, J.; Lin, G.; Xiang, T.; Shao, L.; and Hoi, S. C. 2021. Deep learning for person re-identification: A survey and outlook. *IEEE transactions on pattern analysis and machine intelligence*, 44(6): 2872–2893.

Zhang, Z.; Wang, Y.; Liu, S.; Xiao, B.; and Durrani, T. S. 2021. Cross-domain person re-identification using heterogeneous convolutional network. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(3): 1160–1171.

Zhao, Y.; Zhong, Z.; Yang, F.; Luo, Z.; Lin, Y.; Li, S.; and Sebe, N. 2021. Learning to generalize unseen domains via memory-based multi-source meta-learning for person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6277–6286.

Zheng, L.; Shen, L.; Tian, L.; Wang, S.; Wang, J.; and Tian, Q. 2015. Scalable person re-identification: A benchmark. In *Proceedings of the IEEE international conference on computer vision*, 1116–1124.

Zheng, Z.; Zheng, L.; and Yang, Y. 2017. Unlabeled samples generated by gan improve the person re-identification baseline in vitro. In *Proceedings of the IEEE international conference on computer vision*, 3754–3762.

Zhou, K.; Liu, Z.; Qiao, Y.; Xiang, T.; and Loy, C. C. 2022. Domain generalization: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 45(4): 4396–4415.

Zhou, K.; Yang, Y.; Hospedales, T.; and Xiang, T. 2020. Learning to generate novel domains for domain generalization. In *Computer vision–ECCV 2020: 16th European conference, glasgow, UK, August 23–28, 2020, proceedings, part XVI 16*, 561–578. Springer.

Zhou, K.; Yang, Y.; Qiao, Y.; and Xiang, T. 2021. Domain generalization with mixstyle. *arXiv preprint arXiv:2104.02008*.

Zhuang, W.; Wen, Y.; Zhang, X.; Gan, X.; Yin, D.; Zhou, D.; Zhang, S.; and Yi, S. 2020a. Performance optimization of federated person re-identification via benchmark analysis. In *Proceedings of the 28th ACM international conference on multimedia*, 955–963.

Zhuang, Z.; Wei, L.; Xie, L.; Zhang, T.; Zhang, H.; Wu, H.; Ai, H.; and Tian, Q. 2020b. Rethinking the distribution gap of person re-identification with camera-based batch normalization. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII 16*, 140–157. Springer.