

# Optimizing Performance of Federated Person Re-identification: Benchmarking and Analysis

WEIMING ZHUANG\*, S-Lab, Nanyang Technological University, Singapore

XIN GAN, Nanyang Technological University, Singapore

YONGGANG WEN, Nanyang Technological University, Singapore

SHUAI ZHANG, SenseTime Research, China

The increasingly stringent data privacy regulations limit the development of person re-identification (ReID) because person ReID training requires centralizing an enormous amount of data that contains sensitive personal information. To address this problem, we introduce federated person re-identification (*FedReID*) – implementing federated learning, an emerging distributed training method, to person ReID. *FedReID* preserves data privacy by aggregating model updates, instead of raw data, from clients to a central server. Furthermore, we optimize the performance of *FedReID* under statistical heterogeneity via benchmark analysis. We first construct a benchmark with an enhanced algorithm, two architectures, and nine person ReID datasets with large variances to simulate the real-world statistical heterogeneity. The benchmark results present insights and bottlenecks of *FedReID* under statistical heterogeneity, including challenges in convergence and poor performance on datasets with large volumes. Based on these insights, we propose three optimization approaches: (1) We adopt *knowledge distillation* to facilitate the convergence of *FedReID* by better transferring knowledge from clients to the server; (2) We introduce *client clustering* to improve the performance of large datasets by aggregating clients with similar data distributions; (3) We propose *cosine distance weight* to elevate performance by dynamically updating the weights for aggregation depending on how well models are trained in clients. Extensive experiments demonstrate that these approaches achieve satisfying convergence with much better performance on all datasets. We believe that *FedReID* will shed light on implementing and optimizing federated learning on more computer vision applications.

CCS Concepts: • **Computing methodologies** → **Distributed algorithms**; *Object identification*; *Matching*; • **Information systems** → **Top-k retrieval in databases**.

Additional Key Words and Phrases: federated learning, person re-identification

## ACM Reference Format:

Weiming Zhuang, Yonggang Wen, Xuesen Zhang, Xin Gan, and Shuai Zhang. 2018. Optimizing Performance of Federated Person Re-identification: Benchmarking and Analysis. In . ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/1122445.1122456>

## 1 INTRODUCTION

Person re-identification (ReID) aims to match the same person who appeared in disjoint camera views. It has received considerable attention because of its wide applications in business and public security, such as customer trajectory analysis and criminal investigation [22]. Person ReID has achieved outstanding performances [4, 45, 54], attributing to the advances of deep neural networks (DNN) [13, 21].

\*Also with SenseTime Research.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2018 Association for Computing Machinery.

Manuscript submitted to ACM

However, the increasing concerns of data privacy protection limit the development of person ReID [8]. DNN-based approaches are data-hungry, which relies on centralizing a sizable amount of data to achieve high performance [58]. Training images of person ReID contains sensitive personal information, which could reveal the identity and location of individuals. Centralizing these images impose potential privacy leakage risks. Hence, it is crucial to navigate the development of person ReID under the premise of privacy protection.

Federated learning (FL), an emerging distributed training technique, has empowered many applications with privacy-preserving mechanisms [18], such as healthcare applications [5, 41] and consumer products [36, 37]. FL preserves data privacy by training models collectively with decentralized clients. These clients, instead of transferring raw data, only transfer training updates to a central server. It reduces privacy leakage risks as raw data are kept locally. Despite the advantages of FL, implementing FL to person ReID and optimizing its performance are largely overlooked; such implementation possibility is only mentioned in [12], but that study does not present dataset or benchmark results.

In this work, we propose Federated Person Re-identification (*FedReID*), a new person ReID training paradigm to enable multimedia researchers to train models with privacy guaranteed. Besides privacy protection, FedReID possesses other advantages: reducing communication overhead of uploading plenty amount of data [35]; adapting models in clients to local scenes; obtaining a holistic model that generalizes in diverse scenarios. A usage example of FedReID is video surveillance across communities or districts, where multiple entities collaborate to learn a generalized model without revealing their private video surveillance data.

However, implementing FL to person ReID is not trivial – statistical heterogeneity is a major challenge of FedReID in real-world scenarios [25]: 1) data is in *non-identical and independent distribution (non-IID)* [56] because data collected from different cameras could have significant discrepancies in resolution, illumination, and angles; 2) data volume is *unbalanced* with varied pedestrian flow in different locations. Although some studies illustrate that non-IID harms the training convergence and model performance in tasks like image classification [56], the impact of statistical heterogeneity on FedReID is not explored before.

This work aims to optimize FedReID under statistical heterogeneity via benchmark analysis. We start by constructing a new benchmark, FedReIDBench, with nine representative ReID datasets and a specially designed algorithm for FedReID (Sec. 3). In the benchmark, a server coordinates nine clients (each containing a dataset) to conduct training on their local data and aggregates training updates iteratively. We then conduct benchmark analysis (Sec. 4), revealing that statistical heterogeneity leads to performance degradation and difficulty in convergence. We end by proposing three performance optimization methods: client clustering (Sec. 5.1) and dynamic weight adjustment (Sec. 5.3) to elevate performance; knowledge distillation (Sec. 5.2) to facilitate convergence. Specifically, *client clustering* groups clients with similar data distribution and aggregate training updates within each group. *Knowledge distillation* uses a public dataset to transfer knowledge from clients to the server more effectively. Besides, *weight adjustment* dynamically updates the weights of clients’ training updates in server aggregation. Extensive experiments demonstrate the effectiveness of the benchmark and the significance of the optimization approaches. We believe that FedReID will shed light on implementing and optimizing federated learning to more computer vision applications.

In summary, we make the following contributions:

- We construct a new benchmark for Federated Person Re-identification (FedReID), simulating real-world scenarios of statistical heterogeneity with 9 representative person ReID datasets.
- We provide useful insights and investigate potential bottlenecks of FedReID by analyzing the benchmark results.

- We propose three performance optimization methods: knowledge distillation to facilitate convergence; client clustering and dynamic weight adjustment to elevate performance.
- We extensively evaluate these optimization methods to demonstrate their effectiveness.

## 2 RELATED WORKS

### 2.1 Person Re-Identification

The objective of person ReID is to retrieve the identity of interest from disjoint camera views. It is an important computer vision task that is widely applied in public security, such as video surveillance [58]. The advances of the deep neural network have greatly improved the performance of person ReID by learning better feature representations, compared to traditional hand-crafted feature operations [29, 31, 46, 52]. Over the years of development, the community has constructed many person ReID datasets [15, 28, 50, 57, 59]. These datasets are collected from various locations with different camera views. The majority of person ReID studies focus on extracting better feature representations by improving the architecture of deep neural networks [22, 54]. They rely on the assumption that data, collected from different cameras in various locations, can be centralized to a central server. However, centralizing plenty of images of individuals raises potential privacy leakage risks. Different from previous approaches, we proposed FedReID – a new training paradigm for ReID to learn ReID models from decentralized data. FedReID mitigates potential privacy leakage issues as data is not transferred to a central server.

### 2.2 Federated Learning

Federated learning (FL) is an emerging distributed training technique that trains models with decentralized clients coordinated by a central server [18].

**Benchmarks** To facilitate the development of FL, researchers have published several benchmarks and datasets: LEAF [3] is the first benchmark for FL research, containing federated datasets for image classification and natural language processing tasks; Streets datasets [34] is a real-world image dataset collected from street cameras for object detection; OARF [17] is a benchmark aims to facilitate a wide range of FL applications, such as trend prediction, recommendation, and sentiment analysis. However, different from these tasks, person ReID is a retrieval task that no existing benchmark contains related datasets. In this work, we construct a new FL benchmark that simulates real-world scenarios of FedReID.

**Algorithm** The most known algorithm for FL is Federated Averaging (FedAvg) [35]. It defines an iterative training process that clients send trained local models to a server and the server sends back the aggregated global model to clients. Benchmarks mentioned above adopt FedAvg as the standard algorithm. However, FedAvg requires all clients to have identical models. It is not suitable for FedReID because clients could have varied classifiers. Therefore, we propose an enhanced algorithm, Federated Partial Averaging (FedPav).

**Statistical Heterogeneity** Statistical heterogeneity – non-IID and unbalanced data – is a major challenge of FL [18, 25]. In traditional distributed training [9, 44], data in multiple nodes of cloud clusters are IID. Data in multiple FL clients, however, could be heterogeneous. To address this challenge, some studies focus on optimizing training in clients [1, 19, 24, 26, 55]; while the recent work [55] requires extra communication by sharing features among clients. On the other hand, some studies optimize the aggregation process in the server [47, 48, 61, 62]. Besides, several studies share voluntary or public data between the server and clients [53, 56]. These methods are validated on small datasets

Table 1. The statistics of nine datasets in our proposed benchmark, FedReIDBench. These datasets have large variances in data volume, decreasing from top to bottom.

Datasets	# Cameras	Train		Test			
		# IDs	# Images	Query		Gallery	
				# IDs	# Images	# IDs	# Images
MSMT17 [50]	15	1,041	32,621	3,060	11,659	3,060	82,161
DukeMTMC-reID [59]	8	702	16,522	702	2,228	1,110	17,611
Market-1501 [57]	6	751	12,936	750	3,368	751	19,732
CUHK03-NP [30]	2	767	7,365	700	1,400	700	5,332
PRID2011 [15]	2	285	3,744	100	100	649	649
CUHK01 [28]	2	485	1,940	486	972	486	972
VIPeR [11]	2	316	632	316	316	316	316
3DPeS [2]	2	93	450	86	246	100	316
iLIDS-VID [49]	2	59	248	60	98	60	130

[3, 7, 20], and thus may not be directly applicable to the challenging scenario of FedReID. In this work, we introduce three optimization methods targeting the statistical heterogeneity of FedReID via in-depth benchmark analysis.

This work is an extension of our previous conference version [63]. The main improvements are as follows: 1) We introduce a new performance optimization method — client clustering; 2) We integrate client clustering with the previously proposed weight adjustment method, achieving the best performance; 3) We conduct more performance evaluations for comparison with the benchmark results and the proposed optimization methods; 4) We provide more comprehensive descriptions for the proposed optimization methods. Despite that another work [51] also studies FedReID after our conference work [63], it focuses more on adapting to unseen domains, while we aim to address the statistical heterogeneity revealed from our benchmark analysis.

### 3 FEDERATED PERSON REID BENCHMARK

This section introduces a new federated learning benchmark for person ReID, FedReIDBench. This benchmark comprises nine representative datasets, two possible implementation architectures, one enhanced algorithm, and several performance evaluation metrics.

#### 3.1 Datasets

We construct the benchmark dataset with nine representative person ReID datasets as shown in Table 1. It contains in total 224,064 images of 17,991 identities. These datasets are collected at multiple locations (or countries) and published by different organizations at different times. They not only vary in the number of images, identities, and camera views, but also differ in image resolution, illumination, and scenes.

The variances in these datasets simulate the statistical heterogeneity in real-world scenarios: the disparity of data volumes represents the unbalanced data problem; the domain discrepancies among datasets represent the non-IID problem. Unlike centralized training where data is IID, statistical heterogeneity makes training even more challenging.

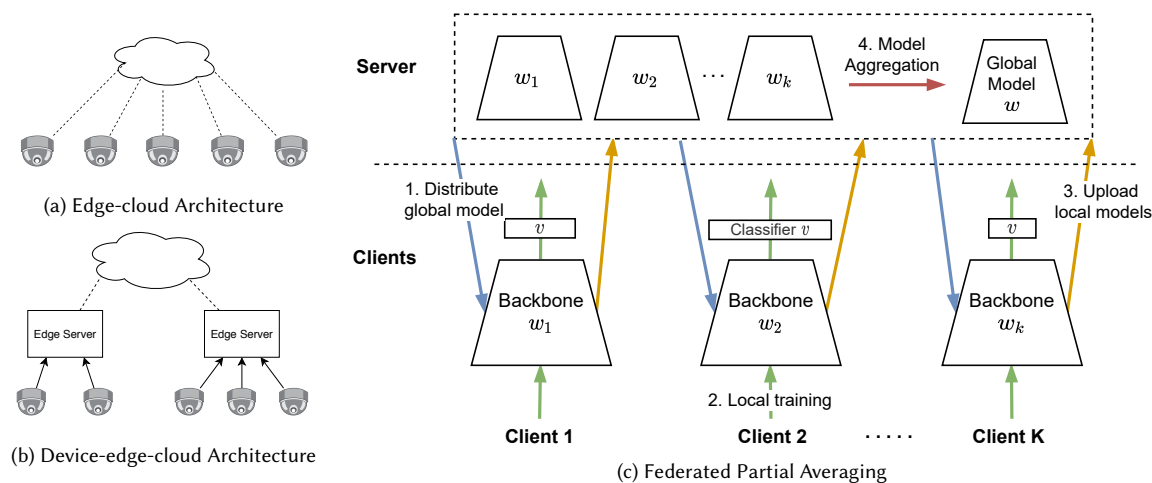


Fig. 1. Illustrations of FedReID architectures and benchmark algorithm Federated Partial Averaging (FedPav). (a) Edge-cloud architecture and (b) device-edge-cloud architecture are two possible architectures for FedReID. (c) FedPav algorithm is suitable for both architectures, regarding cameras as clients for (a) and edge servers as clients for (b). Each training round of FedPav has four steps: 1) A central server distributes a global model to selected clients; 2) Clients conduct training with their local data; 3) Clients upload trained models; 4) The server obtains a new global model by aggregating these models.

### 3.2 Architectures

Figure 1a and Figure 1b illustrate two architectures for possible implementation scenarios of FedReID: edge-cloud architecture and device-edge-cloud architecture. In both architectures, the cloud represents the central server connecting to multiple edges.

**Edge-cloud Architecture:** In this architecture, cameras are the edges that directly connect with the server to conduct federated learning. The server coordinates these cameras to train models with locally collected images. This architecture significantly reduces privacy leakage risks as the data always stays at the edges. However, deployment of this architecture requires cameras to have enough computation power and storage capability. A real-world application of this architecture would be video surveillance for a community with multiple cameras on different streets.

**Device-edge-cloud Architecture:** This is a three-layer hierarchical architecture. Edge servers are in the middle layer. On the one hand, they construct local training datasets by gathering images from multiple camera views, which is similar to how datasets in the benchmark are collected. On the other hand, edge servers collaboratively perform federated learning with their local datasets under the coordination of the server. A good illustration of this architecture would be multiple communities collaborate to learn person ReID models, where each community has an edge server collecting data from multiple cameras.

### 3.3 Algorithm

The standard federated learning algorithm Federated Averaging [35] is not suitable for FedReID because it requires identical model structures in all clients. The model structure of the benchmark is ID-discriminative embedding (IDE) [58], which is a common baseline for DNN-based person ReID. This model structure consists of a backbone and a classifier: the backbone is ResNet-50 [13] in our FedReIDBench; the classifier is a linear layer whose dimension depends

**Algorithm 1:** Federated Partial Averaging (FedPav)

---

**Input:** Total data volume  $n$ , client  $k$ 's data volume  $n_k$ , local epoch  $E$ , batch size  $B$   
**Output:** Global model  $w^T$ , local models of each client  $w_k^T$

---

```

1 Server:
2   initialize  $w^0$ ;
3   for each round  $t = 0$  to  $T-1$  do
4      $S_t \leftarrow$  (random set of  $K$  clients selected from  $N$  clients);            $\triangleright$  Client selection
5     for each client  $k \in S_t$  in parallel do
6        $w_k^{t+1} \leftarrow$  Client( $w^t, k, t$ );                                      $\triangleright$  Distribution
7        $w^{t+1} \leftarrow \sum_{k \in C_t} \frac{n_k}{n} w_k^{t+1}$ ;                                $\triangleright$  Server aggregation
8   return best global model  $w^T$ , best client  $k$ 's local model  $w_k^T$ ;
9 Client ( $w, k, t$ ):
10   $v \leftarrow$  (initialize classifier  $v$  if  $t == 0$ , otherwise retrieve it from local store);
11  for each epoch  $e = 0$  to  $E-1$  do
12    for each batch  $b \in$  batches of local data in size  $B$  do
13       $(w, v) \leftarrow (w, v) - \eta \nabla \mathcal{L}((w, v); b)$ ;            $\triangleright (w, v)$  concatenation of two model structures
14    store  $v$ ;
15  return  $w$ ;                                                          $\triangleright$  Upload local model

```

---

on the number of identities of clients. Since the number of identities could vary among clients, their classifiers could differ in clients. Hence, we adopt the enhanced algorithm for FedReID: Federated Partial Averaging [63].

Federated Partial Averaging (FedPav) allows models in clients to be only partially identical. For FedReID, FedPav enables clients to use the same backbone but different identity classifiers for federated learning, as shown in Figure 1c. The training process is similar to FedAvg except that clients only transfer the identical part of models to the central server for aggregation.

Algorithm 1 summarizes FedPav. We aim to obtain a holistic global model and personalized local models for clients at the end of the training. Each training round  $t$  of FedPav contains four steps: 1) *Distribution*: the central server chooses a fraction ( $K$  out of  $N$ ) of clients for current round of training and distributes the global model  $w^t$  to these clients; 2) *Local training*: each client  $k$  initializes the backbone  $w_k^t$  using the global model parameters and train the model with local dataset for  $E$  local epochs with  $B$  batch size; 3) *Upload*: each client  $k$  uploads the trained backbone  $w_k^{t+1}$  to the server; 4) *Aggregation*: the server generates a new global model  $w^{t+1}$  by aggregating updates from clients with weighted average. The training stops after iterating these four steps for  $T$  rounds. After training, we use the global model  $w$  to evaluate convergence and generalization, and use local models  $w_k$  to evaluate how well models adapt to local scenarios.

### 3.4 Performance Evaluation Metrics

We evaluate FedReID in two aspects: accuracy and communication cost.

**Accuracy** Cumulative Matching Characteristics (CMC) curve and mean Average Precision (mAP) [58] are standard person ReID evaluation metrics. Given an image as a query, person ReID matches it in a gallery of images based on similarity. CMC measures the probability that the query identity is in the top- $k$  most similar matched gallery images. We consider  $k = \{1, 5, 10\}$  in the benchmark, representing the rank-1 accuracy, rank-5 accuracy, and rank-10 accuracy. Besides, we also report the mAP of all queries.

**Communication Cost** Since federated learning requires iterative communication between a server and multiple clients, we also consider the communication costs. The total communication cost is  $T \times 2 \times M$ , where  $T$  is the number of communication rounds and  $M$  is the transmission message size (model size).  $2 \times M$  is the communication cost of each round, considering both uploading and downloading from clients.

### 3.5 Reference Implementation

To facilitate ease-of-use and reproducibility, we open-source referenced implementation in GitHub <sup>1</sup>. It includes data preprocessing, proposed algorithm, and optimization methods. We plan to integrate it to EasyFL [60] in the future. Besides, we provide the experimental settings as follows:

**Learning Rate** The initialized learning rates were different for the identity classifier and the backbone: 0.05 for the identity classifier; 0.005 for the backbone. The learning rate schedulers of both are the same with step size 40 and gamma 0.1. Besides, the learning rate for the server fine-tuning in knowledge distillation is 0.0005.

**Optimizer** We use Stochastic Gradient Decent (SGD) as the optimizer. The optimizer is set with weight decay  $5e-4$  and momentum 0.9.

**Federated Learning Settings** The default settings of federated learning algorithms are as follows: batch size  $B = 32$ , local epoch  $E = 1$ , and total training rounds  $T = 300$ .

## 4 BENCHMARK ANALYSIS

In this section, we present the results of extensive experiments on the benchmark. We investigate the performance of two architectures, the impact of different federated settings, and the impact of statistical heterogeneity.

We initialize the backbone with ResNet-50[13] parameters pre-trained on ImageNet [10]. For hyperparameters, we use batch size  $B = 32$  and local epoch  $E = 1$  to train  $T = 300$  communication rounds by default.

### 4.1 Edge-cloud Architecture

In the edge-cloud architecture, each camera is a client. Since each person ReID dataset contains data from several camera views, we simulate FedReID in this architecture by assigning data of the same camera view to one client. As a dataset is divided into several clients by camera views, we term it *federated-by-camera* scenario.

To understand FedReID performance in the federated-by-camera scenario, we compare it with two other settings: 1) *Federated-by-identity* scenario: we divide one dataset into partitions for multiple clients, where each client includes one partition that contains an equal number of identities. The number of clients equals the number of camera views. 2) *Centralized training*: training with data merged from multiple cameras, which can be considered as the upper bound. For example, Market-1501 dataset [57] contains six camera views with 751 identities. In the federated-by-identity scenario, we divide it into six clients, where each client includes 125 non-overlapping identities. The centralized training means training with the Market-1501 dataset.

Table 2 presents the comparisons of global models of different settings on two datasets: CUHK03-NP dataset [30] and Market-1501 dataset [57]. Compared with the federated-by-identity scenario or centralized training, the federated-by-camera scenario performs much worse. It indicates that learning from only one camera view is infeasible to obtain a generalized model in person ReID, where the evaluation is based on images from multiple camera views. Hence, even though industrial cameras have enough computation and storage capacity to support edge-cloud architecture, the

<sup>1</sup><https://github.com/cap-ntu/FedReID>

Table 2. Performance comparison of federated-by-camera scenario, federated-by-identity scenario, and centralized training on CUHK03-NP and Market-1501 datasets. The federated-by-camera scenario achieves the worst performance, indicating that edge-cloud architecture could be inadequate for FedReID.

Dataset	# Clients	Settings	Rank-1	Rank-5	Rank-10	mAP
CUHK03-NP	2	Federated-by-camera	11.21	19.14	25.71	11.11
		Federated-by-identity	<b>51.71</b>	<b>69.50</b>	<b>76.79</b>	<b>47.39</b>
		Centralized Training	49.29	68.86	76.57	44.52
Market-1501	6	Federated-by-camera	61.13	74.88	80.55	36.57
		Federated-by-identity	85.69	93.44	95.81	66.36
		Centralized Training	<b>88.93</b>	<b>95.34</b>	<b>96.88</b>	<b>72.62</b>

device-edge-cloud architecture could be more adequate for FedReID because each client learns cross-camera knowledge. All the other experiments in the paper are conducted based on the device-edge-cloud architecture.

## 4.2 Device-Edge-Cloud Architecture

In the device-edge-cloud architecture, edge servers collect data from multiple cameras and conduct FedReID with a central server. Since each of the benchmark datasets consists of data from multiple camera views, we simulate this scenario with nine clients – each client contains one unique dataset of the benchmark datasets. In all experiments, we choose nine clients to participate in training.

Under this architecture, We consider two types of models produced from FedReID training: 1) *Local model*: the specialized models trained after  $E$  local epochs in clients before uploading to the server in each training round. 2) *Global model*: the generalized model obtained in the server by aggregating models uploaded from clients.

To understand the performance of FedReID, we compare global and local models with the other two models: 1) *Standalone training*: the models trained in clients with their own dataset (without participating in federated learning); 2) *Centralized training*: the model trained using the combination of all benchmark datasets, simulating conventional person ReID training that centralizes datasets. Centralized training can be treated as the upper bound of FedReID. While FedReID is meaningful for a client only when the performance of global or local models is better than standalone training.

**4.2.1 Impact of Federated Settings.** We first investigate the performance of FedReID (the global model) using the FedPav algorithm under different federated settings, including batch size  $B$  and local epochs  $E$ .

Batch size reflects the trade-off between computation power consumption and model accuracy. With the same local epochs, a larger batch size reduces computation time because the training can better take advantage of the parallelism provided by the client hardware. (Computation is fully utilized as long as  $B$  is large enough.) Figure 2a compares the rank-1 accuracy of FedPav using different batch sizes  $B = \{32, 64, 128\}$ , under the setting that local epochs  $E = 1$  and communication rounds  $T = 300$ . Smaller batch size generally achieves better performance in most datasets, while consumes higher computation.

Local epochs reflect the trade-off between the communication cost and model accuracy. The total training epochs  $E_{total}$  can be calculated with  $E_{total} = T \times E$ , where  $T$  is the communication rounds and  $E$  is the number of local epochs. By fixing the total training epochs for a fair comparison, smaller  $E$  means larger communication rounds  $T$ , requiring higher communication costs. Besides, we compare the rank-1 accuracy of different numbers of local epochs in Figure 2b.






Fig. 2. Performance comparison of (a) different batch sizes and (b) different local epochs. Batch size  $B = 32$  and local epoch  $E = 1$  achieves better overall performance than other settings. We run a total of 300 epochs for these experiments.

Despite that  $E = 5$  performs worse than  $E = 10$  in several datasets, smaller numbers of local epochs  $E$  generally result in better performance. The smallest number of local epoch  $E = 1$  achieves much better performance than  $E = 5$  and  $E = 10$  in all datasets, while it requires the highest communication cost, indicating the trade-off between communication costs and model accuracy.

**4.2.2 Impact of Statistical Heterogeneity.** The statistical heterogeneity hinders the convergence and performance of FedReID. Specifically, non-IID causes difficulty in convergence and both non-IID and unbalanced data limits the performance of FedReID.

Figure 4 shows that FedPav does not converge well as the accuracy (of the global model) fluctuates throughout training. We argue that it is mainly due to non-IID data of nine clients. As datasets in clients have domain discrepancy (e.g. illumination, resolution, scenes, etc.), aggregating them simply by weighted average leads to unstable and unpredictable results. As a result, it causes difficulty in selecting a representative global model for other scenarios. We report the accuracy by averaging the three best global models throughout training, evaluated every ten rounds.

Furthermore, Table 3 compares the performance of the global and local models obtained from FedReID with standalone and centralized training. The results are two-fold: On the one hand, standalone training outperforms both the global and local models in large datasets such as DukeMTMC-reID [59] and CUHK03-NP [30]; On the other hand, both the global and local models outperform standalone training in small datasets such as VIPeR [11] and 3DPeS [2], and even outperforming centralized training in iLIDS-VID dataset [49]. These results indicate that although clients with larger datasets do not benefit from FedReID, the ones with smaller datasets gain significant improvement. We interpret the results from two perspectives: 1) For clients with large datasets, they dominate in server aggregation as the weights for aggregation are positively correlated with data volumes, causing less gain from others; 2) For clients with small datasets, they learn from other clients more effectively because their models are not well trained.

Another observation from Table 3 is that local models outperform the global model in all datasets. As the global model is produced by aggregating local models, we argue that non-IID data causes performance degradation in the server aggregation. Better aggregation methods can be considered to better transfer knowledge from local models to the global model.

## 5 PERFORMANCE OPTIMIZATION

In this section, we first propose three methods to address the problems caused by statistical heterogeneity: client clustering, knowledge distillation, and dynamic weight adjustment. Then, we present experimental results of these optimization methods, compared with standalone training and the benchmark results.

Table 3. Rank-1 accuracy comparison of the global model and local models obtained from FedReID, standalone training, and centralized training. FedReID effectively improves the performance on small datasets. However, it performs worse than standalone training on large datasets due to statistical heterogeneity. We run the experiment with  $B = 32$  and  $E = 1$ .

Methods	MSMT17	DukeMTMC	Market	CUHK03	PRID2011	CUHK01	VIPeR	3DPeS	iLIDS-VID
Centralized Training	54.6	84.2	91.7	64.0	80.0	89.7	65.5	82.1	80.6
Standalone Training	<b>49.6</b>	<b>80.1</b>	<b>88.9</b>	<b>49.3</b>	<b>55.0</b>	69.0	27.5	65.4	52.0
Global Model	41.0	74.3	83.4	31.7	37.7	73.4	48.1	69.2	79.9
Local Model	48.3	78.1	83.6	39.5	50.7	<b>80.7</b>	<b>52.0</b>	<b>80.6</b>	<b>84.7</b>

### 5.1 Client Clustering

To tackle the performance degradation caused by non-IID data in server aggregation of all clients, we propose to aggregate clients with similar data distributions. As discussed in Section 4.2.2, local models outperform the global model in all datasets. The global model is obtained by aggregating local models, so the performance drop mainly sources from the aggregation of clients with diverse data distributions. To tackle this problem, we propose *client clustering* to split clients into several groups based on their data distributions and aggregate models within each group in the server.

Figure 3a depicts the process of client clustering with the following steps: (1) We extract features  $f_k$  from one batch data (32 samples) of a public person ReID dataset<sup>2</sup> using the trained model  $w_k$  from client  $k$ . (2) We adopt a clustering algorithm to cluster these features into multiple groups. (3) We aggregate models of clients within each group, obtaining a global model in each group. (4) We use the global model of each group to update local models of clients within that group for the next training round. In Figure 3a, we cluster clients into two groups: one group contains clients {1, 4} and another one contains clients {2, 3, 5}, based on their features  $f$ . Then, we aggregate  $w_1$  and  $w_4$  to obtain global model  $w_{c1}$ , and aggregate  $w_2$ ,  $w_3$ , and  $w_5$  to obtain  $w_{c2}$ . At the start of next training round, we update local models of clients {1, 4} with  $w_{c1}$  and local models of clients {2, 3, 5} with  $w_{c2}$ . Client clustering obtains multiple global models after training, so we focus on evaluating personalized local models  $w_k$  of each client  $k$ .

In this way, we use the features as a proxy to measure the similarity of data distributions among clients. The intuition behind client clustering is that the clients clustered into the same group share more similar data distributions. The choice of the clustering algorithm is important for the overall performance of this method. We utilize a hierarchical clustering algorithm, FINCH [39], to cluster clients based on similarities of features extracted from their models. Regarding each client as a cluster at the start, we group the clients that are first neighbors; Two clients are first neighbors if their features have the shortest distance (cosine similarity) or they share the same first neighbor. FINCH merges first neighbors in each clustering step. In our scenario, since nine clients would be merged into one cluster after two to three clustering steps, we only cluster for one step per communication round. As a result, the server would have two to three clusters, where each cluster contains two to seven clients. FINCH is able to deliver good clustering results without prior knowledge of the targeted number of clusters.

### 5.2 Knowledge Distillation

Besides client clustering, we adopt *knowledge distillation* (KD) to elevate performance and improve the convergence of FedReID. Since local models outperform the global model, it suggests that local models contain more knowledge than the global model – simple server aggregation could not effectively aggregate knowledge from local models. Knowledge

<sup>2</sup>The public person ReID dataset is shareable among the server and clients. This dataset can be unlabeled.

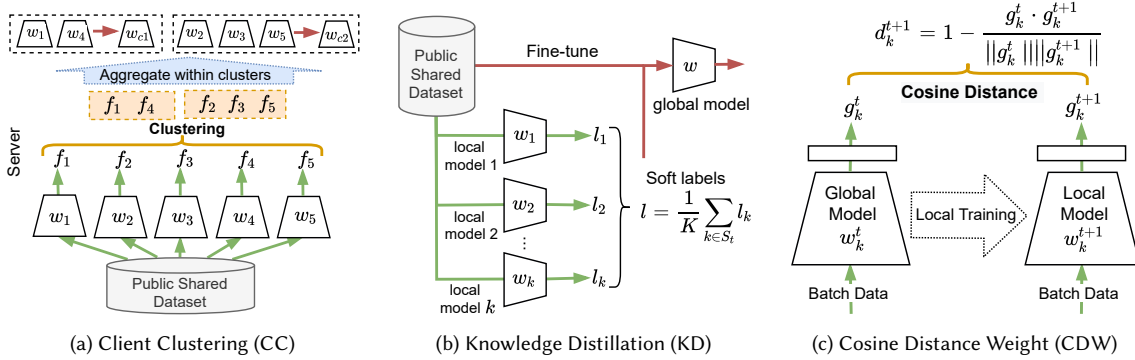


Fig. 3. Illustrations of proposed performance optimization methods: (a) client clustering, (b) knowledge distillation, and (c) cosine distance weight.

distillation is a method proposed by Hinton et. al. [14] to transfer knowledge from a teacher model to a student model, where the teacher model contains more knowledge than the student model. We adopt knowledge distillation to better transfer knowledge from local models to the global model, regarding clients as teachers and the server as the student.

After clients finish local training and upload models, we apply knowledge distillation with a public shared dataset  $\mathcal{D}_{shared}$  in the server. Figure 3b illustrates the additional steps required from knowledge distillation: (1) The server uses each trained model  $w_k$  of client  $k$  to generate soft labels<sup>3</sup>  $l_k$  using samples of  $\mathcal{D}_{shared}$ . These soft labels represent the knowledge of clients' models. (2) Apart from model aggregation, the server aggregates these soft labels with  $l = \frac{1}{K} \sum_{k \in S_t} l_k$ . (3) The server fine-tunes the global model with  $\mathcal{D}_{shared}$  and corresponding labels  $l$  to learn the distilled knowledge.

### 5.3 Weight Adjustment

In addition to tackling the performance degradation caused by non-IID data, we propose to dynamically update the weights for aggregation to curb the adverse effect of unbalanced data. As discussed in Section 4.2.2, the weights of server aggregation are inappropriate. The formula for server aggregation [35, 63] is  $w^{t+1} = \sum_{k \in S_t} \frac{n_k}{n} w_k^{t+1}$ , where  $n$  is the total data volume and  $n_k$  is the data volume of client  $k$ . The weights of local models depend on the data volume of clients – larger datasets lead to larger weights. Since data volumes have large discrepancies among datasets, large datasets dominate in the server aggregation. For example, the weight of the largest datasets (MSMT17 [50] dataset) is around 40%, while the weight of smallest dataset (iLIDS-VID [49]) is only 0.3%. Models from smaller datasets are almost negligible in aggregation. Such unbalanced data volume hampers clients with large datasets to effectively learn from others. Hence, we introduce a novel weight adjustment method to obtain more suitable weights for weighted average in aggregation.

**Cosine Distance Weight** We introduce *cosine distance weight* (CDW) to substitute the weight of data volumes. CDW adjust the weights for aggregation dynamically in each round, based on how well models are trained in clients. It is measured by changes in features extracted from models before and after training. Such changes are calculated by cosine distance. Particularly, in each training round, client  $k$  downloads and trains on the global model  $w_k^t$  from the server to obtain a new local model  $w_k^{t+1}$ . Figure 3c demonstrates our method to calculate the new weight with

<sup>3</sup>These labels are termed soft labels as they are the predicted labels, not the actual labels, of the dataset.

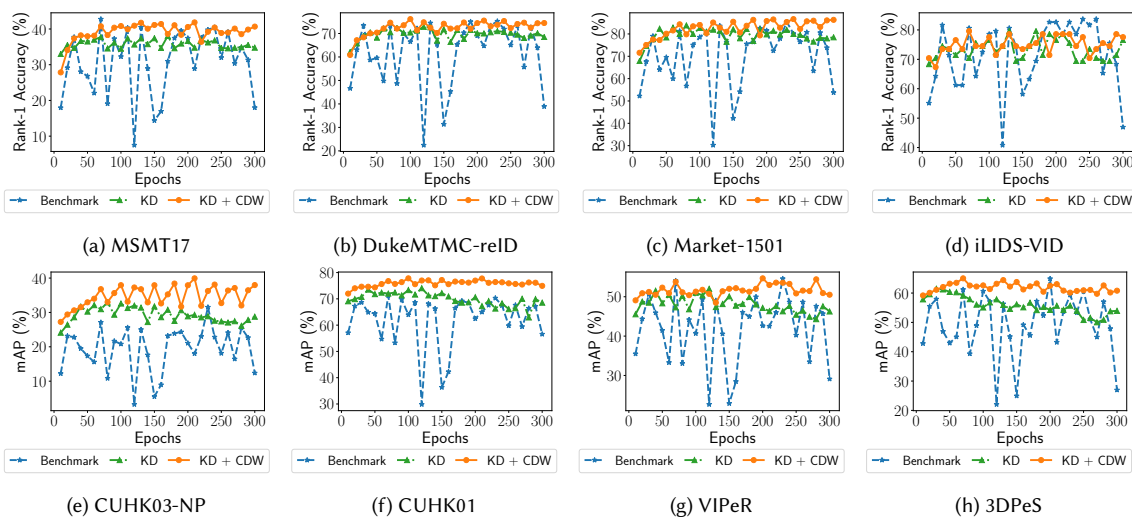


Fig. 4. Convergence and performance (rank-1 accuracy and mAP) comparison of the benchmark results, knowledge distillation (KD), and combination of KD and cosine distance weight (CDW). KD effectively facilitates the convergence of FedReID. Besides, the combination of KD and CDW not only facilitates convergence but also effectively improves performance. These experiments are run with batch size  $B = 32$  and local epoch  $E = 1$ .

$w_k^t$  and  $w_k^{t+1}$ , with following steps: 1) Client  $k$  extracts logits  $g_k^t$  with a random batch data  $\mathcal{D}_{batch}$  using  $(w_k^t, v_k^t)$ <sup>4</sup>. 2) Client  $k$  obtains new local model  $(w_k^{t+1}, v_k^{t+1})$  after local training. 3) Client  $k$  extracts features  $g_k^{t+1}$  with  $\mathcal{D}_{batch}$  using  $(w_k^{t+1}, v_k^{t+1})$ . 4) We calculate the cosine distance of these two logits  $g_k^t$  and  $g_k^{t+1}$ , with following formula:

$$d_k^{t+1} = 1 - \frac{g_k^t \cdot g_k^{t+1}}{\|g_k^t\| \|g_k^{t+1}\|}, \quad (1)$$

where the cosine distance  $d_k^{t+1}$  of each client  $k$  is pushed to the server. The server uses the formula below to obtain the new weight:

$$p_k^{t+1} = \frac{d_k^{t+1}}{\sum_{k \in S_t} d_k^{t+1}}, \quad (2)$$

where the server uses  $p_k^{t+1}$  to replace  $\frac{n_k}{n}$  as the new weight for aggregation. The intuition of CDW is that clients whose local trainings are more effective should contribute more to the aggregation. The cosine distance  $d_k$  measures the scale of changes in local training that updates model  $w_k^t$  to  $w_k^{t+1}$ .

#### 5.4 Combinations of Optimization Method

We can achieve even better performance by combinations of these three optimization methods: client clustering (CC), knowledge distillation (KD), and cosine distance weight (CDW). We consider only combining CDW and CC, and CDW and KD.

It is not desirable to combine client clustering and knowledge distillation because they both enhance the server aggregation. On the one hand, knowledge distillation only fine-tunes a single global model, while client clustering

<sup>4</sup> $(w_k^t, v_k^t)$  is the concatenation of global model  $w_k^t$  and local classifier  $v_k^t$

**Algorithm 2:** FedPav with Client Clustering and Cosine Distance Weight

---

**Input:** Total data volume  $n$ , client  $k$ 's data volume  $n_k$ , local epoch  $E$ , batch size  $B$ , shared dataset  $\mathcal{D}_{shared}$   
**Output:** Global model  $w^T$ , local models of each client  $w_k^T$

1 **Server:**  
2 initialize  $w^0$ ;  
3 **for each round**  $t = 0$  to  $T-1$  **do**  
4      $S_t \leftarrow$  (random set of  $K$  clients selected from  $N$  clients); ▷ Client selection  
5     **for each client**  $k \in S_t$  **in parallel do**  
6          $w_c^t \leftarrow$  (global model of the cluster client  $k$  belongs to);  
7          $w_k^{t+1}, d_k \leftarrow$  **Client**( $w_c^t, k, t$ ); ▷ Distribution  
8          $f_k \leftarrow$  (extract features from  $\mathcal{D}_{batch}$  with  $w_k^{t+1}$ );  
9      $C \leftarrow$  (cluster clients  $S_t$  based on features  $f_k$  for  $k \in S_t$  with FINCH);  
10     **for each client**  $k$  **in each cluster**  $c \in C$  **do**  
11          $d \leftarrow \sum_{k \in S_t} d_k$ ;  
12          $w_c^{t+1} \leftarrow \sum_{k \in S_t} \frac{d_k}{d} w_k^{t+1}$ ; ▷ Server aggregation  
13 **return** best global model  $w^T$ , best client  $k$ 's local model  $w_k^T$ ;

14 **Client** ( $w, k, t$ ):  
15  $v \leftarrow$  (initialize classifier  $v$  if  $t == 0$ , otherwise retrieve it from local store);  
16  $(w^t, v^t) \leftarrow (w, v)$ ; ▷ Save a copy of global model before training  
17 **for each local epoch**  $e = 0$  to  $E-1$  **do**  
18     **for**  $b \in$  batches of local data in size  $B$  **do**  
19          $(w, v) \leftarrow (w, v) - \eta \nabla \mathcal{L}((w, v); b)$ ; ▷  $(w, v)$  concatenation of two model structures  
20      $g, g^t \leftarrow$  (extract logits with  $(w, v)$  and  $(w^t, v^t)$  using  $\mathcal{B}_{shared}$ ); ▷  $\mathcal{B}_{shared}$  is a batch of  $\mathcal{D}_{shared}$   
21      $d = 1 - \frac{g^t \cdot g}{\|g^t\| \|g\|}$ ; ▷ Cosine distance  
22     store  $v^t$ ;  
23 **return**  $w^t, d$ ; ▷ Upload

---

contains multiple global models. On the other hand, both methods address the non-IID problem: knowledge distillation aims to further improve the global model, while client clustering tends to elevate the performance of local models. Hence, we do not consider the combination of these two methods.

Since CDW tackles unbalanced data volume, either the combination of CDW and CC or the combination of CDW and KD addresses statistical heterogeneity with non-IID and unbalanced data problems. To combine CDW with CC or KD, we just need to replace the original weights in the server aggregation process with the new weights. As client clustering has no single global model, combining it with CDW aims to achieve better local models; As knowledge distillation further fine-tunes the global model, combining it with CDW aims to achieve a better global model. We summarize these two combinations in Algorithm 2 and Algorithm 3.

## 5.5 Evaluation

We present the empirical evaluation of these performance optimization approaches, compared with the benchmark and standalone training. By default, we conduct these experiments with batch size  $B = 32$  and local epoch  $E = 1$ . For both client clustering and knowledge distillation, we adopt an additional unlabelled dataset, CUHK02 [27] dataset. This

**Algorithm 3:** FedPav with Knowledge Distillation and Cosine Distance Weight

---

**Input:** Total data volume  $n$ , client  $k$ 's data volume  $n_k$ , local epoch  $E$ , batch size  $B$ , shared dataset  $\mathcal{D}_{shared}$   
**Output:** Global model  $w^T$ , local models of each client  $w_k^T$

1 **Server:**  
2 initialize  $w^0$ ;  
3 **for each round**  $t = 0$  to  $T-1$  **do**  
4      $S_t \leftarrow$  (random set of  $K$  clients selected from  $N$  clients); ▷ Client selection  
5     **for each client**  $k \in S_t$  **in parallel do**  
6          $w_k^{t+1}, d_k \leftarrow$  **Client**( $w^t, k, t$ ); ▷ Distribution  
7          $\ell_k \leftarrow$  (predict soft labels with  $w_k^{t+1}$  and  $\mathcal{D}_{shared}$ );  
8          $d \leftarrow \sum_{k \in S_t} d_k$ ;  
9          $w^{t+1} \leftarrow \sum_{k \in C_t} \frac{d_k}{d} w_k^{t+1}$ ; ▷ Server aggregation  
10          $\ell \leftarrow \frac{1}{K} \sum_{k \in S_t} \ell_k$ ;  
11          $w^{t+1} \leftarrow$  (fine-tune  $w^{t+1}$  with  $\mathcal{D}_{shared}$  and soft labels  $\ell$ );  
12 **return** best global model  $w^T$ , best client  $k$ 's local model  $w_k^T$ ;

13 **Client** ( $w, k, t$ ):  
14  $v \leftarrow$  (initialize classifier  $v$  if  $t == 0$ , otherwise retrieve it from local store);  
15  $(w^t, v^t) \leftarrow (w, v)$ ;  
16 **for each local epoch**  $e = 0$  to  $E-1$  **do**  
17     **for**  $b \in$  batches of local data in size  $B$  **do**  
18          $(w, v) \leftarrow (w, v) - \eta \nabla \mathcal{L}((w, v); b)$ ; ▷  $(w, v)$  concatenation of two model structures  
19  $g, g^t \leftarrow$  (extract logits with  $(w, v)$  and  $(w^t, v^t)$  using  $\mathcal{B}_{shared}$ ); ▷  $\mathcal{B}_{shared}$  is a batch of  $\mathcal{D}_{shared}$   
20  $d = 1 - \frac{g^t \cdot g}{\|g^t\| \|g\|}$ ; ▷ Cosine distance  
21 store  $v$ ;  
22 **return**  $w, d$ ; ▷ Upload

---

dataset is regarded as a public dataset that is shareable among clients and the server. CUHK02 dataset is an extension of the CUHK01 dataset. It includes 7,264 images of 1,816 identities collected from 6 camera views.

We first evaluate the effectiveness of knowledge distillation (KD) and the combination of CDW and KD by monitoring performance changes of global models as training proceeds. Figure 4 shows the performance changes (either rank-1 accuracy or mAP) of KD, the combination of CDW and KD, and the benchmark results on eight datasets. Compared with the benchmark results, training with KD achieves much better convergence; KD can also lead to higher performance, especially when datasets in clients share similar data distributions with the public shared dataset. For example, we use the CUHK02 dataset as the shared dataset, so the accuracy of the global models on both CUHK03-NP and CUHK01 datasets are better than the benchmark results. Moreover, training with the combination of KD and CDW achieves outstanding performance on almost all datasets – better than the benchmark results or training with KD. These results indicate that the combination of KD and CDW is able to obtain the best generalized global model that is transferable to other scenarios.

Next, we evaluate the effectiveness of client clustering (CC), CDW, and the combination of these two methods by comparing the performances of their local models. Table 4 shows the increase in rank-1 accuracy of several methods when compared with standalone training on nine datasets. Although FedNova [48] and FedProx [26] slightly improve

Table 4. Increase in rank-1 accuracy of benchmark results, client clustering (CC), cosine distance weight (CDW), and combination of CC and CDW, when comparing with standalone training. CC effectively improves the performance on larger datasets and CDW effectively elevates the performance on all datasets. Besides, the combination of CC and CDW achieves the best overall performance, especially on the larger datasets. These experiments are run with batch size  $B = 32$  and local epoch  $E = 1$ .

Methods	MSMT17	DukeMTMC	Market	CUHK03-NP	PRID2011	CUHK01	VIPeR	3DPeS	iLIDS-VID
Benchmark	-1.3	-2.0	-5.4	-9.8	-4.3	+11.6	+24.5	+15.2	+32.7
FedNova [48]	-2.1	-2.8	-4.4	-14.6	0.0	+9.9	+24.4	+12.6	<b>+35.8</b>
FedProx [26]	-0.1	-1.6	+1.0	-6.4	-1.0	+12.5	+24.1	+7.7	+34.7
CC	+2.4	-1.3	+0.1	<b>+3.9</b>	+6.0	+9.3	+4.1	-1.2	+16.3
CDW	+4.0	+1.3	+1.4	+1.2	+7.3	<b>+13.8</b>	+26.0	<b>+16.3</b>	+30.3
CC & CDW	<b>+4.1</b>	<b>+3.8</b>	<b>+2.0</b>	+2.2	<b>+13.0</b>	+6.1	<b>+28.2</b>	+6.5	+28.6

Table 5. Rank-1 accuracy comparison of our proposed methods (CC & CDW and KD & CDW) with the existing approaches on two new datasets: CAVIAR [6] and GRID [33]. Our trained models outperform the existing methods on both datasets without extra fine-tuning. These results demonstrate the generalization ability of our methods.

Datasets	Existing methods (w/o privacy except [51])						Ours (w/ privacy)		
	DSTML	UMDL	CrossGrad	MLDG	SSDAL	DIMN	Decentralized [51]	CC & CDW	KD & CDW
CAVIAR [6]	28.2	41.6	-	-	-	-	45.6	46.8	53.2
GRID [33]	-	-	9.0	15.8	22.4	29.3	24.2	30.0	36.8

the performance of the smallest dataset (iLIDS-VID), they are incapable to elevate the performance of large datasets, like our benchmark method. We further analyze the results in three folds. Firstly, CC effectively mitigates the drawback of the benchmark, improving the performance on larger datasets such as MSMT17 [50]. This is because the dominance of larger datasets over smaller datasets is reduced as they are clustered into different groups in aggregation. Most of the time, CC creates two clusters: the first one contains clients with PRID2011, CUHK03-NP, VIPeR, 3DPeS, and iLIDS-VID datasets; the second one contains clients with MSMT17, DukeMTMC-reID, Market-1501, and CUHK01 datasets. Secondly, CDW outperforms the standalone training on all datasets. It indicates that CDW effectively addresses the unbalanced data problem, such that all clients are beneficial in federated learning. Thirdly, the combination of CDW and CC further elevates the performance in most datasets. Although such combination produces slight decreases on smaller datasets compared with CDW, it significantly further improves the performance of larger datasets. It increases the motivation of clients with larger datasets to participate in federated learning.

Lastly, we demonstrate the generalization ability of our methods by comparing existing methods on CAVIAR4REID [6] and GRID [33] datasets. Specifically, we compare with unsupervised cross-domain fine-tuning methods: DSTML [16] and UMDL [38], unsupervised generalization methods: CrossGrad [40], MLDS [23], SSDAL [43], and DIMN [42], as well as the recent work [51]. For evaluation on the CAVIAR4REID, we follow [32, 38] to randomly select 36 identities that appeared on two camera views. GRID dataset contains 250 identities from two camera views. For both datasets, we use images of one camera view as the query and another one as the gallery. Table 5 shows that our proposed FedReID with optimizations (CC & CDW and KD & CDW) outperforms all existing methods on rank-1 accuracy on both datasets; KD & CDW achieves especially good performance. Note that we do not fine-tune trained models on these two evaluation datasets. These results further illustrate the significance of our methods.



## 6 CONCLUSION

In this paper, we present Federated Person Re-identification (FedReID), a new paradigm of person ReID training with decentralized data. To investigate the challenges of FedReID, we construct a new benchmark to simulate real-world scenarios. Based on the results and insights from benchmark analysis, we propose three optimization approaches to elevate performance. We propose client clustering and knowledge distillation to address the non-IID problem and introduce cosine distance weight to address the unbalanced data problem. Empirical results demonstrate that the combination of cosine distance weight and client clustering achieves the best local models, and the combination of cosine distance weight and knowledge distillation achieves the best global model, among all methods. In the future, we plan to investigate the system heterogeneity challenges among clients. We also plan to extend FedReID to support unsupervised learning.

## ACKNOWLEDGMENTS

This study is in part supported by the RIE2020 Industry Alignment Fund - Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s); Singapore MOE under its Tier 1 grant call, Reference number RG96/20; Natioanl Research Foundation, Prime Minister’s Office, Singapore under its Sustainable Tropical Data Centre Testbed programme (STDCT), and its Central Gap Fund (“Central Gap” Award No. NRF2020NRF-CG001-027).

## REFERENCES

- [1] Durmus Alp Emre Acar, Yue Zhao, Ramon Matas, Matthew Mattina, Paul Whatmough, and Venkatesh Saligrama. 2020. Federated learning based on dynamic regularization. In *International Conference on Learning Representations*.
- [2] Davide Baltieri, Roberto Vezzani, and Rita Cucchiara. 2011. 3DPeS: 3D People Dataset for Surveillance and Forensics. In *Proceedings of the 2011 Joint ACM Workshop on Human Gesture and Behavior Understanding (Scottsdale, Arizona, USA) (J-HGBU '11)*. Association for Computing Machinery, New York, NY, USA, 59–64. <https://doi.org/10.1145/2072572.2072590>
- [3] Sebastian Caldas, Peter Wu, Tian Li, Jakub Konecný, H. Brendan McMahan, Virginia Smith, and Ameet Talwalkar. 2018. LEAF: A Benchmark for Federated Settings. *CoRR* abs/1812.01097 (2018). arXiv:1812.01097 <http://arxiv.org/abs/1812.01097>
- [4] Tianlong Chen, Shaojin Ding, Jingyi Xie, Ye Yuan, Wuyang Chen, Yang Yang, Zhou Ren, and Zhangyang Wang. 2019. Abd-net: Attentive but diverse person re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 8351–8361.
- [5] Yiqiang Chen, Xin Qin, Jindong Wang, Chaohui Yu, and Wen Gao. 2020. Fedhealth: A federated transfer learning framework for wearable healthcare. *IEEE Intelligent Systems* (2020).
- [6] D. S. Cheng, M. Cristani, M. Stoppa, L. Bazzani, and V. Murino. 2011. Custom Pictorial Structures for Re-identification. In *British Machine Vision Conference (BMVC)*.
- [7] Gregory Cohen, Saeed Afshar, Jonathan Tapson, and Andre Van Schaik. 2017. EMNIST: Extending MNIST to handwritten letters. In *2017 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2921–2926.
- [8] Bart Custers, Alan M. Sears, Francien Dechesne, Iliana Georgieva, Tommaso Tani, and Simone van der Hof. 2019. *EU Personal Data Protection in Policy and Practice*. Springer.
- [9] Jeffrey Dean, Greg Corrado, Rajat Monga, Kai Chen, Matthieu Devin, Mark Mao, Marc'aurelio Ranzato, Andrew Senior, Paul Tucker, Ke Yang, Quoc V. Le, and Andrew Y. Ng. 2012. Large Scale Distributed Deep Networks. In *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Eds.), Curran Associates, Inc., 1223–1231. <http://papers.nips.cc/paper/4687-large-scale-distributed-deep-networks.pdf>
- [10] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*.
- [11] Douglas Gray and Hai Tao. 2008. Viewpoint Invariant Pedestrian Recognition with an Ensemble of Localized Features. In *In European conference on computer vision*. 262–275.
- [12] Tianshu Hao, Yunyou Huang, Xu Wen, Wanling Gao, Fan Zhang, Chen Zheng, Lei Wang, Hainan Ye, Kai Hwang, Zujie Ren, and Jianfeng Zhan. 2018. Edge AIBench: Towards Comprehensive End-to-end Edge Computing Benchmarking. In *2018 BenchCouncil International Symposium on Benchmarking, Measuring and Optimizing*.
- [13] K. He, X. Zhang, S. Ren, and J. Sun. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 770–778.



- [14] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. Distilling the Knowledge in a Neural Network. In *NIPS Deep Learning and Representation Learning Workshop*. <http://arxiv.org/abs/1503.02531>
- [15] Martin Hirzer, Peter Roth, Csaba Beleznai, and Horst Bischof. 2011. Person Re-Identification by Descriptive and Discriminative Classification. In *Proceedings of the Scandinavian Conference on Image Analysis (SCIA)*. , 91–102.
- [16] Junlin Hu, Jiwen Lu, and Yap-Peng Tan. 2015. Deep transfer metric learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 325–333.
- [17] Sixu Hu, Yuan Li, Xu Liu, Qinbin Li, Zhaomin Wu, and Bingsheng He. 2020. The OARF Benchmark Suite: Characterization and Implications for Federated Learning Systems. *arXiv preprint arXiv:2006.07856* (2020).
- [18] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. 2019. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977* (2019).
- [19] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. 2020. Scaffold: Stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*. PMLR, 5132–5143.
- [20] Alex Krizhevsky, Geoffrey Hinton, et al. 2009. Learning multiple layers of features from tiny images. (2009).
- [21] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* 25 (2012), 1097–1105.
- [22] Qingming Leng, Mang Ye, and Qi Tian. 2019. A survey of open-world person re-identification. *IEEE Transactions on Circuits and Systems for Video Technology* 30, 4 (2019), 1092–1108.
- [23] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. 2018. Learning to generalize: Meta-learning for domain generalization. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [24] Qinbin Li, Bingsheng He, and Dawn Song. 2021. Model-Contrastive Federated Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10713–10722.
- [25] Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. 2020. Federated Learning: Challenges, Methods, and Future Directions. *IEEE Signal Processing Magazine* 37 (2020), 50–60.
- [26] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. 2020. Federated Optimization in Heterogeneous Networks. In *Proceedings of Machine Learning and Systems 2020*. 429–450.
- [27] W. Li and X. Wang. 2013. Locally Aligned Feature Transforms across Views. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*. 3594–3601.
- [28] Wei Li, Rui Zhao, and Xiaogang Wang. 2012. Human Reidentification with Transferred Metric Learning. In *Computer Vision – ACCV 2012*. 31–44.
- [29] W. Li, R. Zhao, T. Xiao, and X. Wang. 2014. DeepReID: Deep Filter Pairing Neural Network for Person Re-identification. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*. 152–159.
- [30] Wei Li, Rui Zhao, Tong Xiao, and Xiaogang Wang. 2014. DeepReID: Deep Filter Pairing Neural Network for Person Re-Identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [31] Jiawei Liu, Zheng-Jun Zha, Qi Tian, Dong Liu, Ting Yao, Qiang Ling, and Tao Mei. 2016. Multi-Scale Triplet CNN for Person Re-Identification. In *Proceedings of the 24th ACM International Conference on Multimedia (Amsterdam, The Netherlands) (MM '16)*. Association for Computing Machinery, New York, NY, USA, 192–196. <https://doi.org/10.1145/2964284.2967209>
- [32] Xiao Liu, Mingli Song, Dacheng Tao, Xingchen Zhou, Chun Chen, and Jiajun Bu. 2014. Semi-supervised coupled dictionary learning for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3550–3557.
- [33] Chen Change Loy, Chunxiao Liu, and Shaogang Gong. 2013. Person re-identification by manifold ranking. In *2013 IEEE International Conference on Image Processing*. IEEE, 3567–3571.
- [34] Jiahuan Luo, Xueyang Wu, Yun Luo, Anbu Huang, Yunfeng Huang, Yang Liu, and Qiang Yang. 2019. Real-World Image Datasets for Federated Learning. *arXiv:1910.11089* [cs.CV]
- [35] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017, 20-22 April 2017, Fort Lauderdale, FL, USA (Proceedings of Machine Learning Research, Vol. 54)*, Aarti Singh and Xiaojin (Jerry) Zhu (Eds.). PMLR, 1273–1282. <http://proceedings.mlr.press/v54/mcmahan17a.html>
- [36] Khalil Muhammad, Qinqin Wang, Diarmuid O’Reilly-Morgan, Elias Tragos, Barry Smyth, Neil Hurley, James Geraci, and Aonghus Lawlor. 2020. Fedfast: Going beyond average for faster training of federated recommender systems. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1234–1242.
- [37] Chaoyue Niu, Fan Wu, Shaojie Tang, Lifeng Hua, Rongfei Jia, Chengfei Lv, Zhihua Wu, and Guihai Chen. 2020. Billion-scale federated learning on mobile clients: A submodel design with tunable privacy. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*. 1–14.
- [38] Peixi Peng, Tao Xiang, Yaowei Wang, Massimiliano Pontil, Shaogang Gong, Tiejun Huang, and Yonghong Tian. 2016. Unsupervised cross-dataset transfer learning for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1306–1315.
- [39] Saquib Sarfraz, Vivek Sharma, and Rainer Stiefelhausen. 2019. Efficient parameter-free clustering using first neighbor relations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 8934–8943.

- [40] Shiv Shankar, Vihari Piratla, Soumen Chakrabarti, Siddhartha Chaudhuri, Preethi Jyothi, and Sunita Sarawagi. 2018. Generalizing across domains via cross-gradient training. *arXiv preprint arXiv:1804.10745* (2018).
- [41] Micah J Sheller, G Anthony Reina, Brandon Edwards, Jason Martin, and Spyridon Bakas. 2018. Multi-institutional deep learning modeling without sharing patient data: A feasibility study on brain tumor segmentation. In *International MICCAI Brainlesion Workshop*. Springer, 92–104.
- [42] Jifei Song, Yongxin Yang, Yi-Zhe Song, Tao Xiang, and Timothy M Hospedales. 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.
- [43] Chi Su, Shiliang Zhang, Junliang Xing, Wen Gao, and Qi Tian. 2016. Deep attributes driven multi-camera person re-identification. In *European conference on computer vision*. Springer, 475–491.
- [44] P. Sun, Y. Wen, R. Han, W. Feng, and S. Yan. 2019. GradientFlow: Optimizing Network Performance for Large-Scale Distributed DNN Training. *IEEE Transactions on Big Data* (2019), 1–1.
- [45] Yifan Sun, Liang Zheng, Yi Yang, Qi Tian, and Shengjin Wang. 2018. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In *Proceedings of the European Conference on Computer Vision (ECCV)*. 480–496.
- [46] Guanshuo Wang, Yufeng Yuan, Xiong Chen, Jiwei Li, and Xi Zhou. 2018. Learning Discriminative Features with Multiple Granularities for Person Re-Identification. In *Proceedings of the 26th ACM International Conference on Multimedia (Seoul, Republic of Korea) (MM '18)*. Association for Computing Machinery, New York, NY, USA, 274–282. <https://doi.org/10.1145/3240508.3240552>
- [47] Hongyi Wang, Mikhail Urochkin, Yuekai Sun, Dimitris Papailiopoulos, and Yasaman Khazaeni. 2020. Federated Learning with Matched Averaging. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=BkluqlSFDS>
- [48] Jianyu Wang, Qinghua Liu, Hao Liang, Gauri Joshi, and H Vincent Poor. 2020. Tackling the objective inconsistency problem in heterogeneous federated optimization. *arXiv preprint arXiv:2007.07481* (2020).
- [49] Taiqing Wang, Shaogang Gong, Xiatian Zhu, and Shengjin Wang. 2014. Person Re-identification by Video Ranking. In *Computer Vision – ECCV 2014*, David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars (Eds.). Springer International Publishing, Cham, 688–703.
- [50] Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. 2018. Person Transfer GAN to Bridge Domain Gap for Person Re-identification. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 79–88.
- [51] Guile Wu and Shaogang Gong. 2021. Decentralised Learning from Independent Multi-Domain Labels for Person Re-Identification. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 4 (May 2021), 2898–2906. <https://ojs.aaai.org/index.php/AAAI/article/view/16396>
- [52] Jiwei Yang, Xu Shen, Xinmei Tian, Houqiang Li, Jianqiang Huang, and Xian-Sheng Hua. 2018. Local Convolutional Neural Networks for Person Re-Identification. In *Proceedings of the 26th ACM International Conference on Multimedia (Seoul, Republic of Korea) (MM '18)*. Association for Computing Machinery, New York, NY, USA, 1074–1082. <https://doi.org/10.1145/3240508.3240645>
- [53] Xin Yao, Tianchi Huang, Rui-Xiao Zhang, Ruiyu Li, and Lifeng Sun. 2019. Federated Learning with Unbiased Gradient Aggregation and Controllable Meta Updating. In *Proceedings of NIPS Federated Learning for Data Privacy and Confidentiality Workshop*.
- [54] Mang Ye, Jianbing Shen, Gaojie Lin, Tao Xiang, Ling Shao, and Steven CH Hoi. 2021. Deep learning for person re-identification: A survey and outlook. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021).
- [55] Lin Zhang, Yong Luo, Yan Bai, Bo Du, and Ling-Yu Duan. 2021. Federated Learning for Non-IID Data via Unified Feature Learning and Optimization Objective Alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 4420–4428.
- [56] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. 2018. Federated Learning with Non-IID Data. *CoRR abs/1806.00582* (2018). [arXiv:1806.00582](http://arxiv.org/abs/1806.00582) <http://arxiv.org/abs/1806.00582>
- [57] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. 2015. Scalable Person Re-identification: A Benchmark. *2015 IEEE International Conference on Computer Vision (ICCV)* (2015), 1116–1124.
- [58] Liang Zheng, Yi Yang, and Alexander G. Hauptmann. 2016. Person Re-identification: Past, Present and Future. *CoRR abs/1610.02984* (2016). [arXiv:1610.02984](http://arxiv.org/abs/1610.02984) <http://arxiv.org/abs/1610.02984>
- [59] Zhedong Zheng, Liang Zheng, and Yi Yang. 2017. Unlabeled Samples Generated by GAN Improve the Person Re-identification Baseline in vitro. In *Proceedings of the IEEE International Conference on Computer Vision*.
- [60] Weiming Zhuang, Xin Gan, Yonggang Wen, and Shuai Zhang. 2022. EasyFL: A Low-code Federated Learning Platform For Dummies. *IEEE Internet of Things Journal* (2022), 1–1. <https://doi.org/10.1109/JIOT.2022.3143842>
- [61] Weiming Zhuang, Xin Gan, Yonggang Wen, Shuai Zhang, and Shuai Yi. 2021. Collaborative Unsupervised Visual Representation Learning From Decentralized Data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 4912–4921.
- [62] Weiming Zhuang, Yonggang Wen, and Shuai Zhang. 2022. Divergence-aware Federated Self-Supervised Learning. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=oVE1z8NINe>
- [63] Weiming Zhuang, Yonggang Wen, Xuesen Zhang, Xin Gan, Daiying Yin, Dongzhan Zhou, Shuai Zhang, and Shuai Yi. 2020. Performance Optimization of Federated Person Re-identification via Benchmark Analysis. In *Proceedings of the 28th ACM International Conference on Multimedia*. 955–963.