

Decentralised Learning from Independent Multi-Domain Labels for Person Re-Identification

Guile Wu, Shaogang Gong
Queen Mary University of London
guile.wu@qmul.ac.uk, s.gong@qmul.ac.uk

Abstract

Deep learning has been successful for many computer vision tasks due to the availability of shared and centralised large-scale training data. However, increasing awareness of privacy concerns poses new challenges to deep learning, especially for human subject related recognition such as person re-identification (Re-ID). In this work, we solve the Re-ID problem by decentralised learning from non-shared private training data distributed at multiple user sites of independent multi-domain label spaces. We propose a novel paradigm called Federated Person Re-Identification (FedReID) to construct a generalisable global model (a central server) by simultaneously learning with multiple privacy-preserved local models (local clients). Specifically, each local client receives global model updates from the server and trains a local model using its local data independent from all the other clients. Then, the central server aggregates transferrable local model updates to construct a generalisable global feature embedding model without accessing local data so to preserve local privacy. This client-server collaborative learning process is iteratively performed under privacy control, enabling FedReID to realise decentralised learning without sharing distributed data nor collecting any centralised data. Extensive experiments on ten Re-ID benchmarks show that FedReID achieves compelling generalisation performance beyond any locally trained models without using shared training data, whilst inherently protects the privacy of each local client. This is uniquely advantageous over contemporary Re-ID methods.

Introduction

In recent years, deep neural network learning has achieved incredible success in many computer vision tasks. However, it relies heavily upon the assumption that a large volume of data can be collected from source domains and stored on a centralised database for model training. Despite the current significant focus on centralised data centres to facilitate big data machine learning drawing from shared data collections, the world is moving increasingly towards localised and private distributed data analysis at-the-edge. This differs inherently from the current assumption of ever-increasing availability of centralised data and poses new challenges to deep learning, especially for human subject related recognition such as person re-identification (Gong et al. 2014).

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

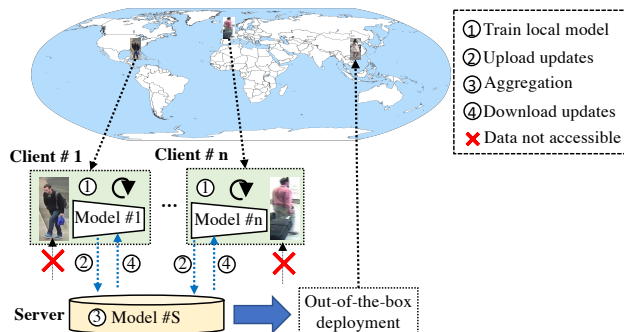


Figure 1: An illustration of FedReID decentralised learning. A client refers to a user site with private local training data, while a server refers to a centralised global model without any training data. Each client uses private local data to train a local model, while a server aggregates local model updates to construct a global model without accessing local training data. This client-server collaborative learning is iterative to yield a global feature representation for out-of-the-box deployment with privacy protection.

Person re-identification (Re-ID) on urban streets at city-wide scales is useful in smart city design (e.g. population flow management) and for public safety (e.g. find a missing person) (Liu, Chang, and Shen 2020; Dong, Gong, and Zhu 2019; Wu, Zhu, and Gong 2020, 2019b). Existing Re-ID methods follow either supervised learning by collecting large-scale datasets for model training (Chen et al. 2020; Liu, Chang, and Shen 2020) or unsupervised learning by assembling both labelled source domain data for model initialisation and unlabelled target domain data for model fine-tuning (Wang and Zhang 2020; Yang et al. 2019b). Although these methods have achieved promising results, they are based on a *centralised learning* paradigm, which is inherently flawed when source datasets cannot be shared in a centralised training protocol due to privacy protection. This requires a new Re-ID paradigm for learning a generalisable global model with *distributed collections of non-shared data from independent multi-domain label spaces*.

In this work, we propose a fundamentally novel paradigm called Federated Person Re-Identification (FedReID) for de-

centralised model learning from distributed non-sharing data of independent label spaces. We construct a generalisable global Re-ID model (a centralised server) by distributed collaborative learning of multiple local models (localised and private clients) without sharing local training data nor collecting any centralised data. As illustrated in Fig. 1, different cities around the world can play the roles of localised clients. Each client receives global model updates from the central server and trains a local model using its own set of private non-shared data. Then, the central server aggregates local model updates to construct a generalisable model without accessing local data. This client-server collaborative learning process is iteratively performed, enabling FedReID to learn a generalisable global model from decentralised data with privacy protection. For deployment, the generalisable global Re-ID model from the server can be deployed directly without using additional centralised data for fine-tuning.

Our *contributions* are: (1) For the first time, we introduce decentralised model learning from distributed non-sharing data of independent multi-domain labels for person Re-ID. This study can potentially benefit other computer vision tasks that require decentralised model learning on distributed non-sharing data with privacy protection. (2) We propose a new paradigm called Federated Person Re-Identification (FedReID). Our approach explores the principle of federated learning (Konečný et al. 2016), but is uniquely designed for decentralised Re-ID by reformulating the iterative client-server collaboration mechanism. In each local client, in addition to a local client model which consists of a feature embedding network for visual feature extraction and a mapping network for classification, we further use a local expert to regularise the training process of the local client model. (3) Extensive experiments on 10 Re-ID benchmarks show that FedReID can both protect local data privacy and achieve compelling generalisation performance, which is uniquely advantageous over contemporary Re-ID methods that assume shared centralised training data without privacy protection.

Related Work

Person Re-Identification. Learning generic feature representations is attractive for real-world Re-ID applications. Conventional supervised Re-ID (Chen et al. 2020; Wu, Zhu, and Gong 2019a) relies heavily on centralised labelled data in target domains, whilst cross-domain unsupervised Re-ID (Wang and Zhang 2020; Yang et al. 2019b) relies on the availability of centralised labelled data from source domains for model initialisation and unlabelled data from target domains for model fine-tuning, so they are both impractical for out-of-the-box deployments. More importantly, the centralised learning paradigm may not be feasible in practice when training data cannot be shared to a centralised training process due to privacy restrictions. Recently, domain generalised Re-ID is proposed to learn a generic feature embedding model. In (Song et al. 2019), Song et al. follow a meta-learning pipeline to optimise a domain-invariant mapping network. In (Jin et al. 2020), a style normalisation module is introduced to filter out style variations. In (Xiao et al. 2016), domain guided dropout is used for domain-specific

knowledge selection. However, these methods still require a centralised training process by assembling a large pool of data from multi-domain datasets. Different from all existing Re-ID methods, our FedReID has a fundamentally new decentralised learning paradigm for optimising a generalised Re-ID model through collaborative learning by communicating knowledge among the central server model and the local client models. Each client learns independently on distributed local private data, while the server uses local model updates to construct a global model without accessing local data nor collecting any centralised data, so FedReID embraces inherently privacy protection.

Federated Learning. Federated learning (Konečný et al. 2016; McMahan et al. 2017; Yang et al. 2019a; Ji et al. 2019) is a recently proposed machine learning technique that allows local users to collaboratively train a centralised model without sharing local data. Conventional federated learning aims at learning a shared model with decentralised data for the same class label space and reducing communication cost. For example, in (McMahan et al. 2017), McMahan et al. introduced Federated Stochastic Gradient Descent (FedSGD) and Federated Average (FedAVG) to iteratively aggregate a shared model by averaging local updates. Our FedReID shares the merit of federated learning but requires a fundamentally different formulation to facilitate the generalisation of a global model for Re-ID. In person Re-ID, each local domain is independent (non-overlapping) from the other domains with totally different person populations (ID space) from different locations/cities, resulting in discrepancies in ID space and context. Thus, we need to learn simultaneously the non-sharing local knowledge in each local client and the latently shared generalisable knowledge in the central server. To this end, in FedReID, each client consists of a feature embedding network for visual feature extraction and a mapping network for classification which is decoupled from the central aggregation process, while the server constructs a generalisable global feature embedding model using updates of local models. Besides, in each local client, we additionally use a local expert to regularise the training process of the local client model to improve the generalisation performance.

Distributed Deep Learning. FedReID differs significantly from distributed deep learning (McClelland and Rumelhart 1989; Dean et al. 2012; Iandola et al. 2016). Distributed deep learning aims at training very large scale deep networks (over billions of parameters) using massive hardware involving tens of thousands of CPU/GPU cores with parallel distributed computation (either model parallelism or data parallelism), with shared large training data. In contrast, FedReID considers the problem of optimising a generalisable global model by asynchronous knowledge aggregation from multi-domain locally learned models without centrally sharing training data.

Private Deep Learning. Private deep learning (Papernot et al. 2017; Wang et al. 2019) aims at constructing privacy-preserving models and preventing the model from inverse attack (Fredrikson, Jha, and Ristenpart 2015; Papernot et al. 2018). A popular solution (Wang et al. 2019) is to use knowledge distillation to transfer private knowledge from multiple teacher ensembles or a cumbersome teacher model to a pub-

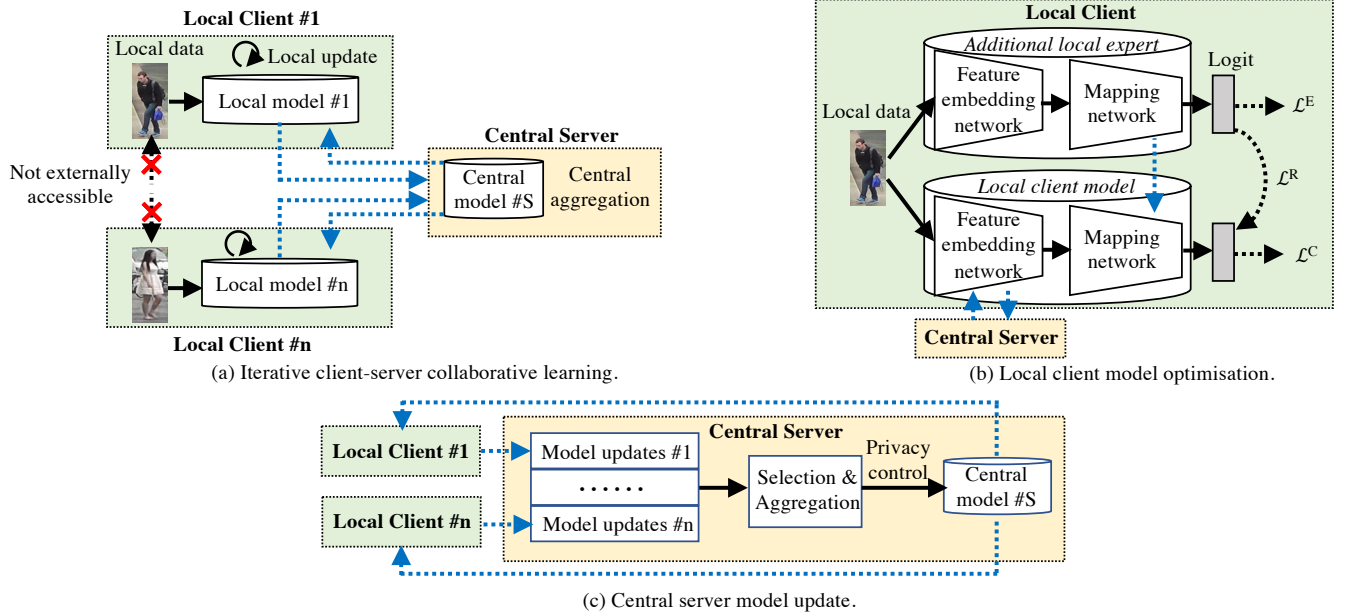


Figure 2: An overview of the proposed Federated Person Re-Identification (FedReID).

lic student model with restricted distillation on training data. In contrast, FedReID does not use any centralised training data (labelled or unlabelled) for model aggregation. Privacy is implemented intrinsically in FedReID by decentralised model training through iterative client-server collaborative learning by asynchronous knowledge aggregation, without central (server) data sharing in model updates.

Methodology

Overview. In this work, we investigate decentralised person Re-ID, a new problem to Re-ID which aims at optimising a generalised model via decentralised learning from independent multi-domain label spaces without assembling local private data. Suppose there are N private datasets captured from different locations that cannot be shared for model training due to privacy protection, *i.e.* there are N localised clients. As shown in Fig. 2(a), each client updates a local model with t_{max} steps separately using its own private data and uploads the model updates to a centralised server. The central server aggregates local model updates to construct a global model and transmits global model updates to each client. This client-server collaborative learning process is iteratively processed, enabling FedReID to learn from decentralised data with privacy protection.

As shown in Fig. 2(b), in each client, there are a local client model and a local expert which are trained together to improve the performance of each client. Specifically, in the i -th client ($i \in N$), suppose there are J_i person images of Z_i identities ($Z_i \leq J_i$) in a local dataset $\mathcal{X}_i = \{x_{i,j}\}_{j=1}^{J_i}$, we construct a feature embedding network $\phi(\omega_{i,t,k}^f)$ to extract feature representations $\mathcal{V}_i = \{v_{i,j}\}_{j=1}^{J_i}$ of person images:

$$v = \phi(\omega_{i,t,k}^f; x) \quad (1)$$

where $\omega_{i,t,k}^f$ are model parameters the i -th feature embedding network at the t -th local step at the k -th global communication epoch. And then, we employ a mapping network $\delta(\omega_{i,t,k}^c)$ for classification:

$$d = \delta(\omega_{i,t,k}^c; v) \quad (2)$$

where d is the outputted logit, $\omega_{i,t,k}^c$ are model parameters of the i -th mapping network at the t -th local step at the k -th global epoch. Meanwhile, we use an additional local expert ($\{\phi(\omega_{i,t,k}^{Ef}), \delta(\omega_{i,t,k}^{Ec})\}$) which is learned with local knowledge of each client and helps the updated local client model to learn richer knowledge¹. Thus, the optimisation objective \mathcal{L}_i of the i -th client is formulated as:

$$\mathcal{L}_i = \mathcal{L}^C + \mathcal{L}^E + \mathcal{L}^R \quad (3)$$

where \mathcal{L}^C is the identity classification loss of the local client, \mathcal{L}^E is the identity classification loss of the local expert, and \mathcal{L}^R is the local expert regularisation from the local expert to the local client model.

As shown in Fig. 2(c), the central server does not use any centralised data for model optimisation. Instead, it selects and aggregates model updates from local clients to construct a server model $\sigma(\theta_k)$ without accessing local private data, where θ_k are model parameters of the server model at the k -th global epoch.

In deployment, the global feature embedding in the central server is directly used for Re-ID matching with a generic distance metric (*e.g.* $L2$ distance).

¹Since the local expert will not be used for the bidirectional client-server knowledge communication, we use “local model” to refer to a local client model which is used for the client-server collaboration and use “local expert” to refer to the additional local expert in each client.

Client-Server Iterative Updates. An intuitive idea for implementing decentralised learning from multiple user sites is to average multiple trained local models in the parameter space to generate a global model. However, this could lead to an arbitrarily bad model (Goodfellow, Vinyals, and Saxe 2015). Recent research in federated learning (Konečný et al. 2016; McMahan et al. 2017) shows that local client models and a central server model can be iteratively updated for distributed model learning. Suppose the i -th client is optimised using SGD with a learning rate η , then its model parameters $\omega_{i,t+1,k}$ at the $(t+1)$ -th local step are updated as:

$$\omega_{i,t+1,k} = \omega_{i,t,k} - \eta \nabla \mathcal{G}_{i,t+1,k} \quad (4)$$

where $\nabla \mathcal{G}_{i,t+1,k}$ is the set of gradient updates of the i -th client at the $(t+1)$ -th local step at the k -th global epoch. After t_{max} steps for local model updates in the clients, the server randomly selects S -fraction ($S \in [0.0, 1.0]$) local clients N_S (here N_S is the set of selected clients) for the server model parameters θ_k aggregation:

$$\theta_{k+1} = \frac{1}{|S \cdot N|} \sum_{m \in N_S} \omega_{m,t_{max},k} \quad (5)$$

where $1 \leq |S \cdot N| \leq N$, $\omega_{m,t_{max},k}$ are model parameters of the m -th client at the t_{max} -th local step of the k -th global epoch. Then, in turn, each client receives θ_k to update the local client model:

$$\omega_{i,0,k+1} = \theta_{k+1} \quad (6)$$

where $\omega_{i,0,k+1}$ are the model parameters of the i -th client at the initial ($t=0$) step of the k -th global epoch. In this way, the local clients and the server are iteratively updated for k_{max} global epochs, and finally we can obtain a global model in the central server for deployment.

FedReID Client-Server Collaboration. In conventional federated learning, all model parameters in the selected client models (including feature embedding layers and classification layers) are used to update the centralised server model (Eq. (5)). However, in decentralised Re-ID, aggregating all model parameters might lead to performance degradation in both local and global models, because each local dataset is usually captured in different locations where the person ID space and context are non-overlapping. To optimise a centralised model across different domains, we reformulate federated learning to simultaneously consider the non-sharing local knowledge in each client and the latently shared generalisable knowledge in the central server.

Specifically, we decouple $\omega_{i,t,k}^c$ (the mapping network) from the aggregation in Eqs. (5) and (6) to preserve local classification knowledge in each client, and aggregate $\omega_{i,t,k}^f$ (feature embedding network) to construct a generalisable feature embedding model for deployment. Starting the feature embedding network of each local client model from the same initialisation, we can accumulate updates of multiple local feature embedding networks to find wider optima in the parameter space of a global model. Thus, in each local client, Eq. (6) is reformulated as:

$$\{\omega_{i,0,k+1}^f, \omega_{i,0,k+1}^c\} = \{\theta_{k+1}^f, \omega_{i,t_{max},k}^c\} \quad (7)$$

Since local data in different clients are from different domains, $\omega_{i,t,k}^c$ in each client corresponds to classification knowledge for different domains. Thus, in the central server, $\omega_{i,t,k}^c$ does not need to be averaged. Besides, since the feature embedding network of each local client starts from the same initialisation (Eq. (7)), accumulating local updates of feature embedding networks corresponds to aggregating model parameters in the feature embedding space. Therefore, the aggregation process in Eq.(5) can be formulated as:

$$\begin{aligned} \theta_{k+1}^f &= \theta_k^f - \frac{\eta}{|S \cdot N|} \sum_{m \in N_S} \sum_{t=1}^{t_{max}} \nabla \mathcal{G}_{m,t,k}^f \\ &= \frac{1}{|S \cdot N|} \sum_{m \in N_S} (\theta_k^f - \eta \sum_{t=1}^{t_{max}} \nabla \mathcal{G}_{m,t,k}^f) \\ &= \frac{1}{|S \cdot N|} \sum_{m \in N_S} \omega_{m,t_{max},k}^f \end{aligned} \quad (8)$$

where θ_k^f are model parameters of the feature embedding network of the central server model, $\nabla \mathcal{G}_{m,t,k}^f$ is the set of gradient updates of the m -th local feature embedding network at the t -th local step at the k -th global epoch.

Optimisation Objective. In FedReID, the central server does not use any centralised data for model fine-tuning, so its optimisation process is the selection and aggregation process as formulated in Eq. (8). In each local client, as shown in Fig. 2(b) and Eq. (7), the local client model receives global model updates from a central model. Then, we use a cross-entropy loss to learn classification knowledge:

$$\mathcal{L}^C = - \sum_{z=1}^{Z_i} y_z \log \frac{\exp(d_z)}{\sum_{b=1}^{Z_i} \exp(d_b)} \quad (9)$$

where y_z is the ground-truth label and d_z is the logit over a class z . To further improve generalisation of the local client model, we use a local expert to help the local client model to learn richer knowledge via knowledge distillation (Hinton, Vinyals, and Dean 2015), which also potentially facilitates the aggregation in the global model. Specifically, the local expert with model parameters $\{\omega_{i,0,k+1}^{Ef}, \omega_{i,0,k+1}^{Ec}\}$ is initialised with Eq. (10) and also optimised with a cross-entropy loss with Eq. (11):

$$\{\omega_{i,0,k+1}^{Ef}, \omega_{i,0,k+1}^{Ec}\} = \{\omega_{i,t_{max},k}^f, \omega_{i,t_{max},k}^c\} \quad (10)$$

$$\mathcal{L}^E = - \sum_{z=1}^{Z_i} y_z \log \frac{\exp(d'_z)}{\sum_{b=1}^{Z_i} \exp(d'_b)} \quad (11)$$

where d' is the logit computed by Eqs. (1) and (2) with model parameters of the local expert. As shown in Fig. 2(b), Eqs. (7) and (10), the difference between the local client model and the local expert is that the former receives a global model for starting feature embedding network with the same initialisation among clients, while the latter utilises the latest local feature embedding network as an initialisation. Since the local expert is only learned with local data without receiving external updates, it performs better than

the global model on the local domain but worse than the global model on the other domains (likely overfit locally). Thus, we use it as a regularisation to help the updated local client model to learn richer knowledge on a specific local domain. We feed the same batch to the two models but with separately random data augmentation and compute soft probability distributions for the local client model (\mathcal{P}) and the local expert (\mathcal{Q}) as:

$$\mathcal{P}_z = \frac{\exp(d_z/T)}{\sum_{b=1}^{Z_i} \exp(d_b/T)}, \mathcal{Q}_z = \frac{\exp(d'_z/T)}{\sum_{b=1}^{Z_i} \exp(d'_b/T)} \quad (12)$$

where T is a temperature (Hinton, Vinyals, and Dean 2015). The local expert regularisation \mathcal{L}^R is therefore defined as the KL-divergence between \mathcal{P} and \mathcal{Q} :

$$\mathcal{L}^R = T^2 \sum_{z=1}^{Z_i} \mathcal{Q}_z \cdot \log \frac{\mathcal{Q}_z}{\mathcal{P}_z} \quad (13)$$

In summary, each local client is optimised with Eq. (3) using local private data independent from the other clients.

Privacy Protection. In FedReID, local sensitive data are inherently protected by the decentralised learning process. To further protect sensitive data from inverse attack (Fredrikson, Jha, and Ristenpart 2015), we use the white noise (Geyer, Klein, and Nabi 2017) to hide the contributions of each client in Eq. (8):

$$\theta_{k+1}^f = \frac{1}{[S \cdot N]} \sum_{m \in N_S} \omega_{m,t_{max},k}^f + \beta \mathcal{N}(0,1) \quad (14)$$

where $\mathcal{N}(0,1)$ is the white noise matrices with mean 0 and variance 1. $\beta \in [0,1]$ is a scale factor to control the balance between privacy-preserving and Re-ID accuracy. When $\beta=0$, the white noise is removed from the aggregation. Moreover, in the client-server collaboration, we can also hide the collaboration information in Eq. (7) as:

$$\{\omega_{i,0,k+1}^f, \omega_{i,0,k+1}^c\} = \{\theta_{k+1}^f + \beta \mathcal{N}(0,1), \omega_{i,t_{max},k}^c\} \quad (15)$$

Experiments

Datasets. We used four large-scale Re-ID datasets (DukeMTMC-ReID (Zheng, Zheng, and Yang 2017), Market-1501 (Zheng et al. 2015), CUHK03 (Li et al. 2014; Zhong et al. 2017) and MSMT17 (Wei et al. 2018)) as non-shared local datasets in four client sites. Each of the four local clients did not share its training data with other clients nor the server. This is significantly different from existing domain generalised Re-ID, where FedReID is trained with decentralised data, while existing methods are trained with centralised data. The FedReID model was then evaluated on five smaller Re-ID datasets (VIPeR (Gray and Tao 2008), iLIDS (Zheng, Gong, and Xiang 2009), 3DPeS (Baltieri, Vezzani, and Cucchiara 2011), CAVIAR (Cheng et al. 2011) and GRID (Loy, Liu, and Gong 2013)), plus a large-scale Re-ID dataset (CUHK-SYSU person search (Xiao et al. 2017)) as new unseen target domains for out-of-the-box deployment tests. For CUHK-SYSU, we used ground-truth person bounding box annotations for Re-ID test (not person search), of which there are 2900 query persons and each

Types	Datasets	Tr. id	Tr. img	Te. id	Te. img
Local Client Train	Duke	702	16522	-	-
	Market	751	12936	-	-
	CUHK03	767	7365	-	-
	MSMT17	1041	30248	-	-
New Domain Test	VIPeR	-	-	316	632
	iLIDS	-	-	60	120
	3DPeS	-	-	96	192
	CAVIAR	-	-	36	72
	GRID	-	-	125	1025
	Cuhk-Sysu	-	-	2900	8347

Table 1: The Re-ID dataset statistics. 'Tr.': Train; 'Te.': Test; 'id': number of identities; 'img': number of images.

person contains at least one image in the gallery (both query and gallery sets are fixed removing distractors in the variable gallery sets). On smaller Re-ID datasets, we did random half splits to generate 10 training/testing splits. In each split, we randomly selected one image of each test identity as the query while the others as the gallery for evaluation. The dataset statistics are summarised in Table 1. Besides, we used CIFAR-10 (Krizhevsky and Hinton 2009) for federated formulation generalisation analysis on image classification.

Evaluation Metrics. We used Rank-1 (R1) accuracy and mean Average Precision (mAP) for Re-ID performance evaluation. Note that FedReID is designed to learn a generalised model from distributed datasets with privacy protection.

Implementation Details. In our design, the feature embedding network is ResNet-50 (He et al. 2016) (pretrained on ImageNet), while the mapping network consists of two fully connected layers, in which the first fully connected layer follows by a batch normalization layer, a ReLU layer and Dropout. In practice, both global and local models used a multi-head architecture in which each mapping network in each head was corresponding to one client. By default, we set the number of local clients $N=4$, client fraction $S=1.0$, and privacy control parameter $\beta=0$. These hyperparameters can be determined by different application requirements. We empirically set batch size to 32, maximum global communication epochs $k_{max}=100$, maximum local steps $t_{max}=1$, and temperature $T=3$. We used SGD as the optimiser with Nesterov momentum 0.9 and weight decay $5e-4$. The learning rates were set to 0.01 for embedding networks and 0.1 for mapping networks, which decayed by 0.1 every 40 epochs. Our models were implemented with Python(3.6) and PyTorch(0.4), and trained on TESLA V100 GPU (32GB).

Federated Formulation Generalisation Analysis

To analyse the generalisation of our FedReID, we compared FedReID with FedSGD (McMahan et al. 2017), FedAVG (McMahan et al. 2017), and FedATT (Ji et al. 2019). When adapting FedSGD, FedAVG and FedATT for Re-ID, we set the last classification layer with the maximal identity number among local datasets. When adapting FedReID for CIFAR-10, we aggregate all layers (including classifiers, *i.e.* using Eqs. (5) and (6)) in the client-server collaboration because all local data are from the same domain. On

Methods	CIFAR-10	VIPeR	iLIDS
FedSGD	90.72±0.04	41.0	65.2
FedAVG ($S=0.5$)	93.04±0.19	41.5	65.3
FedAVG ($S=1.0$)	93.21±0.08	41.0	65.2
FedATT ($S=0.5$)	92.86±0.12	38.3	65.3
FedATT ($S=1.0$)	92.97±0.21	40.2	61.3
FedReID ($S=0.5$)	93.14±0.14	45.3	68.3
FedReID ($S=1.0$)	93.31±0.11	46.2	69.7
Baseline (Centralised joint)	93.58±0.14	44.6	65.5

Table 2: Evaluating generalisation of federated formulations. Top-1/R1 accuracies are reported. FedSGD means setting $S=1.0$ and $t_{max}=1$.

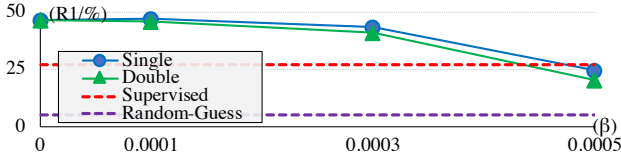


Figure 3: Evaluating the privacy control parameter β on VIPeR. ‘Single’: use β in Eq. (14), ‘Double’: use β in both Eqs. (14) and (15), ‘Random-Guess’: initialise the model with ImageNet pretrained parameters without training.

CIFAR-10, we employed ResNet-32 for experiments and set $t_{max}=5$. As shown in Table 2, FedReID achieves better generalisation than the other federated variants. When training and testing on the same domain for image classification (CIFAR-10), FedReID performs closely to the centralised joint-training, but FedReID only accesses to local model updates with privacy protection whilst joint-training assembles together all local data without privacy concerns. When testing out-of-the-box on unseen new domains for decentralised Re-ID (VIPeR and iLIDS), FedReID performs even better than joint-training, which can be attributed to: (1) FedReID accumulate multiple local model updates to find wider optima in the parameter space of a global model; (2) With the client-server collaboration, FedReID learns softer distribution knowledge. Therefore, FedReID might get better generalisation for an unseen new Re-ID domain while joint-training is the upper bound for testing on a source client.

Privacy Protection Analysis

The inherent defensive ability of FedReID is given by decentralised learning and model aggregation. Besides, β in Eqs. (14) and (15) can further control privacy protection. From Fig. 3, we can see that: (1) R1 accuracy of FedReID gradually decreases when β increases, but FedReID achieves significantly better accuracy than the Random-Guess; (2) Single β protection performs slightly better than double protection, but the double one protects more information; (3) When $\beta=0.0005$, R1 accuracy of FedReID remains close to the overfitting local supervised method, which indicates the compromise of accuracy and privacy.

Settings	Methods	VIPeR	iLIDS
Individuals	Client (Duke)	25.0	56.2
	Client (Market)	26.1	48.0
	Client (CUHK03)	21.6	41.0
	Client (MSMT)	27.3	60.5
Ensembles	Feat-Concatenation	29.4	56.2
	Parameter-Average	19.9	41.8
Decentralised	FedReID	46.2	69.7
Centralised	Baseline (Joint)	44.6	65.5

Table 3: Comparison with individual clients and ensembles. R1 accuracies are reported.

Settings	Methods	VI.	iL.	3D.	CA.	GR.
w/o privacy (Cross-domain fine-tune)	TJAIDL	38.5	-	-	-	-
	DSTML	8.6	33.4	32.5	28.2	-
	UMDL	31.5	49.3	-	41.6	-
	PAUL	45.2	-	-	-	-
w/o privacy (Centralised generalised)	SyRI	43.0	56.5	-	-	-
	SSDAL	43.5	-	-	-	22.4
	MLDG [†]	23.5	53.8	-	-	15.8
	CrossGrad [†]	20.9	49.7	-	-	9.0
	DIMN	51.2	70.2	-	-	29.3
w/ privacy (Decentralised)	FedReID	46.2	69.7	67.0	45.6	24.2

Table 4: Generalised Re-ID performance evaluation on VIPeR, iLIDS, 3DPeS, CAVIAR and GRID. R1 accuracies are reported. [†]: Re-ID domain generalisation results.

Comparison with Individuals and Ensembles

We separately trained the baseline model on four localised datasets as the individuals and used feature concatenation and model parameter average as the ensembles. As shown in Table 3: (1) FedReID significantly outperforms the individuals and the ensembles, which shows that the collaboration between the localised clients and the centralised server facilitates holistic optimisation, enabling FedReID to construct a better generalisable model with privacy protection; (2) Compared with the centralised joint-training baseline, FedReID achieves competitive R1 accuracies, demonstrating its effectiveness; (3) Averaging multiple trained local models in the parameter space leads to an arbitrarily bad model.

Generalised Re-ID Performance Evaluation

FedReID is uniquely designed for protecting local client privacy by learning a generalisable model without centralised sharing of training data, whilst no existing Re-ID methods explicitly integrate privacy protection requirements into their designs. To further show the generalisation of FedReID, we employed out-of-the-box evaluation on non-client unseen domains without training data for fine-tuning.

Results on Smaller Benchmarks. We compared FedReID with: (1) four unsupervised cross-domain fine-tuning methods (TJAIDL (Wang et al. 2018), DSTML (Hu, Lu, and Tan 2015), UMDL (Peng et al. 2016), PAUL (Yang et al. 2019b)), and (2) five unsupervised generalisation methods (SyRI (Bak, Carr, and Lalonde 2018), SSDAL (Su et al.

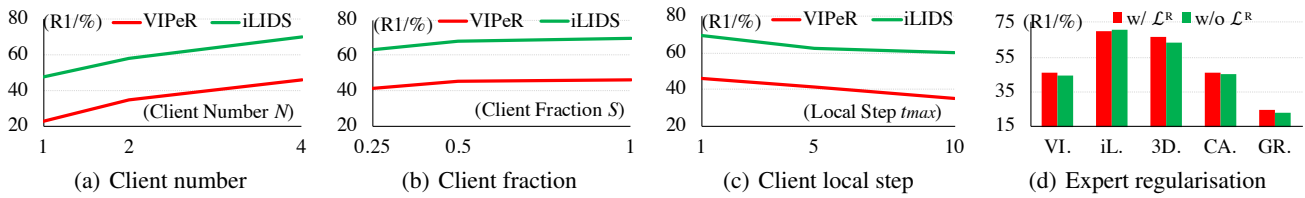


Figure 4: Ablation study on (a) client number, (b) client fraction, (c) local steps and (d) expert regularisation.

Methods	mAP	R1	R5	R10
DSIFT*+Euclidean	41.1	45.9	-	-
BoW*+Cosine	62.5	67.2	-	-
DLDP*	74.0	76.7	-	-
FedReID	80.4	83.4	90.3	92.4
Baseline (Centralised joint)	74.7	77.4	87.2	90.1

Table 5: Evaluation on CUHK-SYSU person search subset for Re-ID. *: Reported results using ground-truth person images and a gallery size of 100 images per query.

Methods	Market		Duke	
	R1	mAP	R1	mAP
Baseline(local sup.)	88.3	71.4	77.3	58.9
FedReID(direct w/o fine-tuning)	80.2	60.1	68.0	52.1
FedReID(sup. w/ local data&dis.)	90.3	76.2	77.4	60.7
DPR(unsup. multi-domain dis.)	61.5	33.5	48.4	29.4
DPR(semi-sup. multi-domain dis.)	63.7	35.4	57.4	36.7

Table 6: Source labelled client tests. On a larger benchmark MSMT17, FedReID (direct) yields 48.4% R1, FedReID (supervised) 64.6% R1, and Baseline 61.7% R1, respectively. sup.: supervised, unsup.: unsupervised, dis.: distillation.

2016), MLDG (Li et al. 2018), CrossGrad (Shankar et al. 2018), DIMN (Song et al. 2019)). As shown in Table 4, FedReID performs competitively against contemporary methods, which shows the effectiveness of the generalised global model for out-of-the-box deployment with privacy protection. For example, FedReID achieves 46.5% R1 on VIPeR and 69.7% R1 on iLIDS, which are the second-best and close to DIMN (Song et al. 2019) which assembles all training data together without data privacy protection. Note that, FedReID is also compatible with other techniques for better generalisation, such as style normalisation (Jin et al. 2020).

Results on Large Benchmark. To further evaluate FedReID on a large-scale target domain, we used the Re-ID subset of CUHK-SYSU person search dataset, which has distinctively different scene context to most other Re-ID datasets above. As shown in Table 5, FedReID achieves competitive performance compared with some unsupervised Re-ID methods (DSIFT (Zhao, Ouyang, and Wang 2013), BoW (Zheng et al. 2015) and DLDP (Schumann, Gong, and Schuchert 2017)) and the centralised join-training baseline, which shows the generalisation of the global model of FedReID for deployment on large-scale Re-ID.

Further Analysis and Discussion

Source Client Tests. As shown in Table 6, although FedReID (direct) only aggregates local model updates *w/o accessing local data*, it still achieves competitive performance, especially when compared with DPR (Wu et al. 2019) which uses local unlabelled data (unsupervised or semi-supervised). We also tested FedReID (supervised) by using local data for model training and a global model as the expert for distillation, which improves FedReID (direct) and outperforms Baseline (local supervised). Here, the improvement in R1 on Duke is small because Duke contains many ambiguous hard negatives and knowledge distillation helps to learn softer distribution for retrieving more positives but is not so helpful for distinguishing hard negatives.

Client Number N . Fig. 4(a) compares central server aggregation with different numbers of local clients, where $N=1, 2$ and 4 denote Market, Market+Duke and Market+Duke+CuHK03+Msmt as clients, respectively. We can see that that collaboration of multi-domain clients is good for learning more generalisable knowledge in the central server.

Client Fraction S . Fig. 4(b) compares FedReID with different client fractions S . We can see that updating with an arbitrary client ($S=0.25$) is inferior to aggregating multiple clients, whilst aggregating all clients ($S=1.0$) performs slightly better than aggregating randomly selected clients ($S=0.5$), but random selection protects more data privacy.

Client Local Step t_{max} . Fig. 4(c) compares FedReID with different client local steps which potentially promote communication efficiency. Overall, the performance of FedReID gradually decreases when t_{max} increases due to the accumulation of biases in each local client.

Local Expert Regularisation. Fig. 4(d) shows the evaluation on the regularisation of local experts. With the expert regularisation, FedReID gets better generalisation overall, which shows that the expert regularisation provides richer knowledge to facilitate the improvement of FedReID.

Conclusion

In this work, we introduced decentralised learning from independent multi-domain label spaces for person Re-ID and proposed a new paradigm called Federated Person Re-Identification (FedReID). We trained multiple local models at different clients using non-shared private local data and aggregated local model updates to construct a global model in a central server. This iterative client-server collaborative learning helps to build a generalisable model for deployment. Extensive experiments on ten Re-ID datasets show the unique advantage of FedReID over contemporary methods.

Acknowledgements

This work is supported by Vision Semantics Limited, Alan Turing Institute Turing Fellowship, and Innovate UK Industrial Challenge Project on Developing and Commercialising Intelligent Video Analytics Solutions for Public Safety (98111-571149), Queen Mary University of London Principal's Scholarship.

References

- Bak, S.; Carr, P.; and Lalonde, J.-F. 2018. Domain adaptation through synthesis for unsupervised person re-identification. In *ECCV*.
- Baltieri, D.; Vezzani, R.; and Cucchiara, R. 2011. 3DPeS: 3D people dataset for surveillance and forensics. In *ACM J-HGBU*.
- Chen, X.; Fu, C.; Zhao, Y.; Zheng, F.; Song, J.; Ji, R.; and Yang, Y. 2020. Saliency-Guided Cascaded Suppression Network for Person Re-Identification. In *CVPR*.
- Cheng, D. S.; Cristani, M.; Stoppa, M.; Bazzani, L.; and Murino, V. 2011. Custom pictorial structures for re-identification. In *BMVC*.
- Dean, J.; Corrado, G.; Monga, R.; Chen, K.; Devin, M.; Mao, M.; Ranzato, M.; Senior, A.; Tucker, P.; Yang, K.; et al. 2012. Large scale distributed deep networks. In *NeurIPS*.
- Dong, Q.; Gong, S.; and Zhu, X. 2019. Person search by text attribute query as zero-shot learning. In *ICCV*.
- Fredrikson, M.; Jha, S.; and Ristenpart, T. 2015. Model inversion attacks that exploit confidence information and basic countermeasures. In *ACM CCS*.
- Geyer, R. C.; Klein, T.; and Nabi, M. 2017. Differentially private federated learning: A client level perspective. In *NeurIPS*.
- Gong, S.; Cristani, M.; Yan, S.; and Loy, C. C. 2014. *Person Re-Identification*. Springer. ISBN 1447162951, 9781447162957.
- Goodfellow, I. J.; Vinyals, O.; and Saxe, A. M. 2015. Qualitatively characterizing neural network optimization problems. In *ICLR*.
- Gray, D.; and Tao, H. 2008. Viewpoint invariant pedestrian recognition with an ensemble of localized features. In *ECCV*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Hu, J.; Lu, J.; and Tan, Y.-P. 2015. Deep transfer metric learning. In *CVPR*.
- Iandola, F. N.; Moskewicz, M. W.; Ashraf, K.; and Keutzer, K. 2016. Firecaffe: near-linear acceleration of deep neural network training on compute clusters. In *CVPR*.
- Ji, S.; Pan, S.; Long, G.; Li, X.; Jiang, J.; and Huang, Z. 2019. Learning private neural language modeling with attentive aggregation. In *IJCNN*.
- Jin, X.; Lan, C.; Zeng, W.; Chen, Z.; and Zhang, L. 2020. Style normalization and restitution for generalizable person re-identification. In *CVPR*.
- Konečný, J.; McMahan, H. B.; Yu, F. X.; Richtárik, P.; Suresh, A. T.; and Bacon, D. 2016. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*.
- Krizhevsky, A.; and Hinton, G. 2009. Learning multiple layers of features from tiny images. Technical report.
- Li, D.; Yang, Y.; Song, Y.-Z.; and Hospedales, T. M. 2018. Learning to generalize: Meta-learning for domain generalization. In *AAAI*.
- Li, W.; Zhao, R.; Xiao, T.; and Wang, X. 2014. Deep-reid: Deep filter pairing neural network for person re-identification. In *CVPR*.
- Liu, C.; Chang, X.; and Shen, Y.-D. 2020. Unity Style Transfer for Person Re-Identification. In *CVPR*.
- Loy, C. C.; Liu, C.; and Gong, S. 2013. Person re-identification by manifold ranking. In *ICIP*.
- McClelland, J. L.; and Rumelhart, D. E. 1989. *Explorations in parallel distributed processing: A handbook of models, programs, and exercises*. MIT press.
- McMahan, H. B.; Moore, E.; Ramage, D.; Hampson, S.; et al. 2017. Communication-efficient learning of deep networks from decentralized data. In *AISTATS*.
- Papernot, N.; Abadi, M.; Erlingsson, U.; Goodfellow, I.; and Talwar, K. 2017. Semi-supervised knowledge transfer for deep learning from private training data. In *ICLR*.
- Papernot, N.; Song, S.; Mironov, I.; Raghunathan, A.; Talwar, K.; and Erlingsson, Ú. 2018. Scalable private learning with pate. In *ICLR*.
- Peng, P.; Xiang, T.; Wang, Y.; Pontil, M.; Gong, S.; Huang, T.; and Tian, Y. 2016. Unsupervised cross-dataset transfer learning for person re-identification. In *CVPR*.
- Schumann, A.; Gong, S.; and Schuchert, T. 2017. Deep learning prototype domains for person re-identification. In *ICIP*.
- Shankar, S.; Piratla, V.; Chakrabarti, S.; Chaudhuri, S.; Jyothi, P.; and Sarawagi, S. 2018. Generalizing across domains via cross-gradient training. In *ICLR*.
- Song, J.; Yang, Y.; Song, Y.-Z.; Xiang, T.; and Hospedales, T. M. 2019. Generalizable Person Re-identification by Domain-Invariant Mapping Network. In *CVPR*.
- Su, C.; Zhang, S.; Xing, J.; Gao, W.; and Tian, Q. 2016. Deep attributes driven multi-camera person re-identification. In *ECCV*.
- Wang, D.; and Zhang, S. 2020. Unsupervised Person Re-identification via Multi-label Classification. In *CVPR*.
- Wang, J.; Bao, W.; Sun, L.; Zhu, X.; Cao, B.; and Philip, S. Y. 2019. Private model compression via knowledge distillation. In *AAAI*.

Wang, J.; Zhu, X.; Gong, S.; and Li, W. 2018. Transferable joint attribute-identity deep learning for unsupervised person re-identification. In *CVPR*.

Wei, L.; Zhang, S.; Gao, W.; and Tian, Q. 2018. Person transfer GAN to bridge domain gap for person re-identification. In *CVPR*.

Wu, A.; Zheng, W.-S.; Guo, X.; and Lai, J.-H. 2019. Distilled person re-identification: Towards a more scalable system. In *CVPR*.

Wu, G.; Zhu, X.; and Gong, S. 2019a. Person re-identification by ranking ensemble representations. In *ICIP*.

Wu, G.; Zhu, X.; and Gong, S. 2019b. Spatio-temporal associative representation for video person re-identification. In *BMVC*.

Wu, G.; Zhu, X.; and Gong, S. 2020. Tracklet self-supervised learning for unsupervised person re-identification. In *AAAI*.

Xiao, T.; Li, H.; Ouyang, W.; and Wang, X. 2016. Learning deep feature representations with domain guided dropout for person re-identification. In *CVPR*.

Xiao, T.; Li, S.; Wang, B.; Lin, L.; and Wang, X. 2017. Joint detection and identification feature learning for person search. In *CVPR*.

Yang, Q.; Liu, Y.; Chen, T.; and Tong, Y. 2019a. Federated machine learning: Concept and applications. *ACM TIST* .

Yang, Q.; Yu, H.-X.; Wu, A.; and Zheng, W.-S. 2019b. Patch-Based Discriminative Feature Learning for Unsupervised Person Re-Identification. In *CVPR*.

Zhao, R.; Ouyang, W.; and Wang, X. 2013. Unsupervised salience learning for person re-identification. In *CVPR*.

Zheng, L.; Shen, L.; Tian, L.; Wang, S.; Wang, J.; and Tian, Q. 2015. Scalable person re-identification: A benchmark. In *ICCV*.

Zheng, W.-S.; Gong, S.; and Xiang, T. 2009. Associating groups of people. In *BMVC*.

Zheng, Z.; Zheng, L.; and Yang, Y. 2017. Unlabeled samples generated by gan improve the person re-identification baseline in vitro. In *ICCV*.

Zhong, Z.; Zheng, L.; Cao, D.; and Li, S. 2017. Re-ranking person re-identification with k-reciprocal encoding. In *CVPR*.