Privacy-Enhancing Person Re-identification Framework – A Dual-Stage Approach

Kajal Kansal, Yongkang Wong, and Mohan Kankanhalli

School of Computing, National University of Singapore

kajal.kansal@nus.edu.sg, yongkang.wong@nus.edu.sg, mohan@comp.nus.edu.sg

Abstract

*In this work, we show that deep learning-based re-identification (Re-ID) models, albeit trained only with a Re-ID objective (*i.e. *if two samples belong to the same identity), encode personally identifiable information (PII) in the learned features that may lead to serious privacy concerns. In cognizance of the modern privacy regulations on protecting PII, we propose a novel dual-stage person Re-ID framework that (1) suppresses the PII from the discriminative features, and (2) introduces a controllable privacy mechanism through differential privacy. The former is achieved with a self-supervised de-identification (De-ID) decoder and an adversarial-identity (Adv-ID) module, whereas the latter mechanism leverages a controllable privacy budget to generate a privacy-protected gallery with a Gaussian noise generator. Furthermore, we introduce the notion of a privacy metric to quantify the privacy leakage in Re-ID features which is not explicitly examined in prior work. We demonstrate the feasibility of our approach in achieving a better trade-off between utility and privacy through rigorous experiments on person Re-ID benchmarks.*

1. Introduction

Person re-identification (Re-ID) is an important task in video surveillance systems as it enables value-added applications, such as person tracking across multiple cameras, forensics, and security [42]. Briefly, the task of person Re-ID aims to match individuals (without identity inference) across multiple cameras that have non-overlapping views. Formally, given a person observed on one camera (*i.e.* a query), a Re-ID system will match that against a restricted set of observed persons from other cameras (*i.e.* gallery) and return the best matching candidates. The key challenges are to deal with visual variations, such as poses, lighting conditions, scale differences, and changing backgrounds, as well as to handle the non-overlapping identities (disjoint identities) between the training dataset and actual deployment.

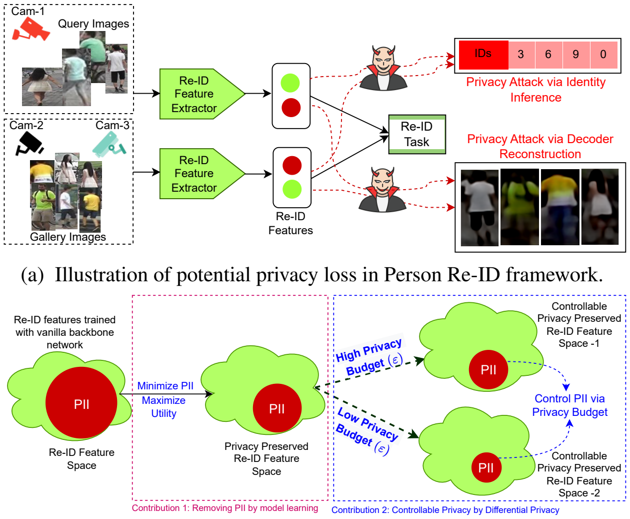
Cam-1
Query Images
Cam-2
Cam-3
Gallery Images
Re-ID Feature Extractor
Re-ID Feature Extractor
Re-ID Features
Re-ID Task
IDs
3
6
9
0
Privacy Attack via Identity Inference
Privacy Attack via Decoder Reconstruction
Re-ID features trained with vanilla backbone network
PII
Re-ID Feature Space
Minimize PII
Maximize Utility
PII
Privacy Preserved Re-ID Feature Space
High Privacy Budget (ε)
Low Privacy Budget (ε)
Controllable Privacy Preserved Re-ID Feature Space -1
PII
Control PII via Privacy Budget
Controllable Privacy Preserved Re-ID Feature Space -2
PII
Contribution 1: Removing PII by model learning
Contribution 2: Controllable Privacy by Differential Privacy

(a) Illustration of potential privacy loss in Person Re-ID framework.

(b) A conceptual illustration of Re-ID feature space obtained from arbitrary trained Re-ID models (left), our privacy-preserving Re-ID model (middle), and with controllable privacy (right).

Figure 1. Overview of privacy leakage in person Re-ID framework (top) and our dual-stage contributions (bottom).

Re-ID models [12, 18, 21, 40] primarily focus on improving the Re-ID utility and overlook the privacy concerns for individuals (*i.e.* person identification). This is because identity inference models cannot be applied directly in practice, as the observed individuals during deployment are not present in the training. Therefore, existing deep learning-based Re-ID models [31, 37, 39] indiscriminately learn the fine-grained appearance features that could be used to infer personally identifiable information (PII), such as identity or other personal attributes (*e.g.* physical appearance, gender, body size, or clothing information). We illustrate the potential privacy threats in Fig. 1a, where a Re-ID feature that encodes fine-grained PII can be exploited for identity inference or reconstructing the observed images. We empirically show the existence of this phenomenon in Section 5.

To avoid the potential of PII being misused, there is an increasing demand for privacy-preserving person Re-ID models that can satisfy the contradictory requirements of simultaneously achieving personal privacy protection and high re-identification accuracy. In practice, in order to enable person Re-ID utility in a surveillance system, the operator needs to define a finite set of detected persons from a particular camera(s) to form a gallery. Motivated by the fact that a learned privacy-preserving model may still leak privacy, *i.e.* perfect protection is often impossible in practice, we note that it is equally important to protect the gallery while matching each query against the gallery.

In this work, we address the privacy leakage in Re-ID based surveillance systems with a dual-stage person Re-ID framework (see Fig. 1b). The proposed framework comprises two components. First, we propose to suppress the PII from an arbitrary Re-ID model via a self-supervised de-identification (De-ID) decoder and an adversarial-identity (Adv-ID) module. Second, we aim to allow users to control the degree of privacy in the gallery with a differential privacy mechanism after learning the privacy-preserving Re-ID features. Here, a user controllable privacy budget is leveraged to generate a privacy-protected gallery with a Gaussian noise generator. In addition to the novel framework, we introduce the notion of privacy metric which can quantify the privacy leakage (*i.e.* PII) in the learned Re-ID models. The key contributions of this work are as follows.

- The proposed privacy-preserving Re-ID model aims to maintain the person Re-ID performance as well as suppress PII via adversarial identity inference tasks. Specifically, we propose to learn privacy-preserved Re-ID features via de-identification (De-ID) and adversarial identity (Adv-ID) information to suppress PII from the learned features.

- We enable the user to control the degree of privacy of a protected gallery and term it *controllable privacy-preserving person Re-ID*. Specifically, we apply differential privacy (DP) on the gallery via a Gaussian noise generator, where the added noise can be controlled via dedicated privacy budget parameters. This approach protects the gallery via added perturbation generated with respective privacy budgets. The protected gallery can be used to match the queries and retrieve the best matching candidates in a privacy-preserving manner.

- We empirically demonstrate the phenomenon of privacy leakage information via person identification with the existing Re-ID models. Further, the trade-offs between user privacy protection and Re-ID utility are shown by re-identification (Re-ID) and identification (ID) performance at different privacy budgets. Extensive experiments show that our approach achieves a better trade-off on competitive Re-ID benchmarks.

2. Related Work

2.1. Person Re-identification (Re-ID)

Person Re-ID targets to match image pairs of a person across non-overlapping camera views [32, 41, 42]. Most of the existing deep learning-based Re-ID models explore fine-grained pedestrian feature descriptions and have shown significant progress [12, 17, 21]. To further increase the Re-ID performance, diverse auxiliary information has been incorporated into these deep Re-ID networks. For example, [34] detected person pose landmarks to obtain human body regions. Lin *et al*. [22] exploited camera ID information to assist inter-image similarity estimation. [23] encoded detailed local descriptors and find the person attributes to improve the Re-ID performance. Barbosa *et al*. [4] demonstrated the use of depth maps to exploit soft-biometric cues. Sun *et al*. [35] utilized part-level features by multiple classifiers, which offer further finer granularity for pedestrian images. Due to their fine-grained nature, these Re-ID models completely ignore the issue of privacy leakage of individuals. Training images of person Re-ID contain PII that could reveal the identity of individuals. Hence, it is crucial to develop a person Re-ID model that avoids potential privacy leakage risks. Here, our main focus is to work towards a privacy-preserved high quality Re-ID model while preventing it from obtaining sensitive visual information that can intrude upon people's privacy.

2.2. Privacy Protection methods

The privacy concern about disclosing personally identifiable information (PII) is data misuse. In relation to data misuse, various legal regulations are introduced, such as the European General Data Protection Regulation [5] and Personal Data Protection Act [8], to protect individual's privacy rights. These privacy laws also stipulate that images of a person are personal data. However, there is no optimal solution to address the image privacy concerns of Person Re-ID. One way to comply with strict data privacy regulations is data anonymization [33] which can allow data to be used/analyzed without compromising the identities of individuals. To address the privacy-related issues, existing anonymization methods use de-identification techniques [3, 19], such as pixelization, blurring, and obfuscation by removing the identity-related cues from data. However, most of these techniques tend to remove the semantic information and compromise the usability of the data. Additionally, identity-irrelevant utilities can shift the identity of a person which can hamper the training process for robust Re-ID models. Here, we propose a novel approach to utilize de-identification techniques in a principled manner that can anonymize the data while ensuring that the anonymization does not negatively affect the Re-ID.

Re-ID data can be easily misused for identity theft, profiling, harassment, and blackmailing and poses a severe threat to individuals' privacy. Some research efforts [9, 29] have been applied towards privacy preserving Re-ID on human faces via blurring. However, they completely ignore other privacy cues that can reveal PII such as individual behavior, location, clothing, body type, *etc*. Moreover, these methods are more focused on preserving Re-ID performance instead of preserving privacy. Hence, privacy-preservation via removing PII is not fully addressed in the literature. Marina *et al*. [30] uses an RGBD camera and Ahmad *et al*. [2] used event cameras to address privacy problems in Re-ID. However, this arrangement is not feasible and scalable in a practical environment. Zhao *et al*. [43] and Cheng *et al*. [7] utilized an encryption strategy and exploited encrypted feature vectors to ensure privacy in Re-ID. However, encrypting huge amounts of visual data is a complex procedure. Zhuang *et al*. [45] uses federated learning (FL) and tries to preserve data privacy by aggregating model updates. Wu *et al*. [38] propose a FedReID model based on a decentralized learning paradigm to construct a global model by simultaneously learning with multiple local models preserved for privacy. However, statistical heterogeneity is a major challenge in the implementation of federated learning-based methods and special efforts are required to design new benchmarks. Differently from these methods, we explore a simple, yet effective mechanism that can remove PII as well as provide a degree of freedom to control privacy while preserving the Re-ID performance.

2.3. Differential Privacy

Differential privacy (DP) is the standard privacy protection approach that offers strong privacy guarantees. DP was proposed by Dwork *et al*. [11] and is to first derive the frequency distribution of the tuples in the input data and then publish a noisy version of the distribution to preserve privacy. Subsequently, it has become the state-of-the-art privacy paradigm for sanitizing statistical databases. DP is utilized extensively in healthcare applications [14, 28] to preserve patient privacy. For example, Meng *et al*. [28] integrates local differential privacy (LDP) and locality-sensitive hashing techniques into the recommendation model to address privacy concerns. LDP is also investigated for geolocation data [36]. He *et al*. [14] secure the privacy of the shared biomedical data via controlling the released information with the help of DP. Chamikara *et al*. [6] applies privacy-calibrated perturbation to biometric data with the help of DP and provides privacy to human faces. Inspired by the success of DP, we use the power of DP to achieve privacy guarantees for Re-ID. However, it is challenging to apply to unstructured and non-aggregated data. Here, our first goal is to study the feasibility of introducing DP in Re-ID by proposing an efficient mechanism. With this, we apply it to the Re-ID gallery database at the deployment stage where the system operator can obtain a degree of freedom to control privacy via adding different levels of the perturbations and can run standard Re-ID algorithms to match the users' queries against the protected database.

3. Problem Definition

This work considers the problem of privacy preserving person Re-ID. Given a training set with N images, $(\boldsymbol{x}_i)_{i=1}^N$, belonging to K identities, the key objective is to train an image encoder $\mathcal{M}_{\texttt{enc}}(\boldsymbol{x}_i)$ to produce a corresponding feature $\boldsymbol{f}_i$ for person Re-ID task. Specifically, $\mathcal{M}_{\texttt{enc}}$ need to satisfy two criteria: (**i**) $\mathcal{D}(\boldsymbol{f}_i, \boldsymbol{f}_j) < \mathcal{D}(\boldsymbol{f}_i, \boldsymbol{f}_k)$ where $\boldsymbol{f}_i$ & $\boldsymbol{f}_j$ belong to the same identity, $\boldsymbol{f}_k$'s identity is different from $\boldsymbol{f}_i$ & $\boldsymbol{f}_j$, and $\mathcal{D}$ is an arbitrary distance function (*e.g*. Euclidean distance), and (**ii**) the identity information is suppressed in $\boldsymbol{f}_i$ such that $\boldsymbol{f}_i$ cannot be used for the person identification task.

To deploy a trained Re-ID model, a set of images from a fixed camera with non-overlapped identities from $(\boldsymbol{x}_i)_{i=1}^N$ is first used to form a gallery database $\mathcal{G}$. Then, any queries from $\mathcal{Q}$, which are detected from a different camera(s) from $\mathcal{G}$, can be used to retrieve the closest matches from $\mathcal{G}$.

4. Proposed Approach

4.1. Overview

In this work, we propose a novel framework that balances the trade-off between utility (Re-ID) and privacy (ID), such that it can learn privacy-preserved Re-ID features for persons. In addition, the proposed framework also provides a degree of freedom to control privacy. As illustrated in Fig. 2, our proposed framework, namely Privacy Enhancing Person Re-ID Network (PEPR-Net), has a Re-ID encoder $\mathcal{M}_{\texttt{enc}}$ to learn Re-ID representations $\boldsymbol{f}_i$, a de-identified decoder $\mathcal{M}_{\texttt{dec}}$ to inject the anonymity information into the learned representation, and an adversarial supervision $\mathcal{M}_{\texttt{adv}}$ to remove the identity information. To provide controllable privacy, we apply differential privacy through a Gaussian noise generator to generate a privacy-protected gallery with dedicated privacy budget parameters.

4.2. Baseline Network

Our backbone network is a feature encoder based on ResNet-50. We embed a fully connected (FC) layer of 128 units over the head of the last FC layer. Then, we utilized the FC layer to extract the Re-ID feature representation, which can be used for describing the visual appearance of the person's image. We train the encoder network using triplet verification loss such that it can minimize the distance between positive pairs ($\boldsymbol{x}_i$ and $\boldsymbol{x}_j$) and maximize the distance between negative pairs ($\boldsymbol{x}_i$ and $\boldsymbol{x}_k$).

The triplet verification loss for the encoder to perform

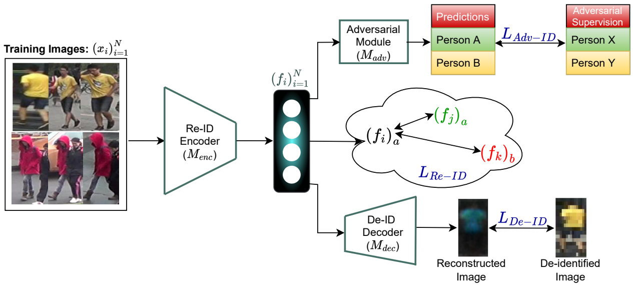

Figure 2. **Proposed framework.** We input $(\boldsymbol{x}_i)_{i=1}^N$ images to the Re-ID encoder ($\mathcal{M}_{\texttt{enc}}$) where N is the number of images and trained with $\mathcal{L}_{\texttt{Re-ID}}$ loss to extract the features ($\boldsymbol{f}_i$). We fed $\boldsymbol{f}_i$ to the De-ID decoder ($\mathcal{M}_{\texttt{dec}}$) to inject De-ID information in $\boldsymbol{f}_i$ through reconstructing the images with the supervision of de-identified images by optimizing $\mathcal{L}_{\texttt{De-ID}}$ loss. $\boldsymbol{f}_i$ features are also fed to the adversarial module ($\mathcal{M}_{\texttt{adv}}$) to remove the identity information with adversarial supervision via $\mathcal{L}_{\texttt{Adv-ID}}$ loss.

Re-ID is given below:

$$\mathcal{L}_{\texttt{Re-ID}} = \frac{1}{\mathcal{T}} \sum_{i=1}^{\mathcal{T}} \Bigg[\Big\| \mathcal{M}_{\texttt{enc}}(\boldsymbol{x}_i) - \mathcal{M}_{\texttt{enc}}(\boldsymbol{x}_j) \Big\|_2^2 - \Big\| \mathcal{M}_{\texttt{enc}}(\boldsymbol{x}_i) - \mathcal{M}_{\texttt{enc}}(\boldsymbol{x}_k) \Big\|_2^2 + \beta \Bigg]_+ \tag{1}$$

where a positive constant β represents the margin, $\boldsymbol{x}_i$ is an anchor query image, x_j is the positive sample (or same identity) for anchor, $\boldsymbol{x}_k$ is the negative sample (or different identity) for anchor. $M_{\texttt{enc}}(\cdot)$ is the output of the encoder, $\mathcal{T}$ is the number of triplets and $[\cdot]_+$ is the Hinge function.

4.3. Privacy preserving Re-ID Model

4.3.1 De-identified Decoder

The feature encoder ($\mathcal{M}_{\texttt{enc}}$) is trained together with a de-identified decoder ($\mathcal{M}_{\texttt{dec}}$) that tries to reconstruct the de-identified image via deidentification. De-identification (De-ID) is a method to remove personal information from data by removing the association between a set of identifying data and the data subject. The de-identified images can be generated with image distortion methods such as pixelization, blurring, and obfuscation. For the de-identified decoder, we use the pixelization [15] technique on the original images to create the de-identified versions and to protect the individual identity. We further provide these deidentified images as target images while supervising the deidentified decoder. The goal of augmenting the network with a de-identified decoder is to reconstruct the de-identified version of the original image by calculating the difference between the reconstructed image and the de-identified original image. Gradients that flow during backpropagation can entangle the deidentified information back into the features, resulting filtering out PII information from Re-ID features while enhancing the feature robustness for only Re-ID. Thus, it helps in circumventing privacy challenges by reconstructing samples with deidentified images.

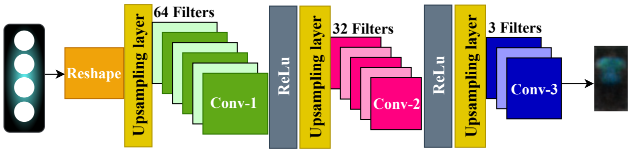

Figure 3. The architecture of the De-ID Decoder network.

The decoder network consists of three convolution layers, three upsampling layers, and two activation layers (as shown in Fig. 3). The output of the feature encoder $\mathcal{M}_{enc}$ is reshaped, upsampled, and given as input to the first convolution layer of the decoder. Each convolution layer is followed by an activation layer where we use ReLu as an activation function and upsampling layer to unpool the samples. After the final convolution layer, the reconstructed image is obtained as an output. The loss function for the de-identified decoder is calculated between the output image and the target deidentified image which is given by $\mathcal{L}_{\texttt{De-ID}}$ as follows:

$$\mathcal{L}_{\texttt{De-ID}} = \sum_{i=1}^{N} \Big\| \hat{\boldsymbol{x}}_i - \mathcal{M}_{\texttt{dec}}\Big(\mathcal{M}_{\texttt{enc}}(\boldsymbol{x}_i)\Big) \Big\|_2^2 \tag{2}$$

where $\boldsymbol{x}_i$ is the original input image, target ground-truth de-identified images are denoted by $\hat{\boldsymbol{x}}_i$, where $\mathcal{M}_{\texttt{enc}}(.)$ and $\mathcal{M}_{\texttt{dec}}(.)$ are the output of the encoder and decoder and N is the number of images.

4.3.2 Adversarial Module

For the adversarial module $\mathcal{M}_{\texttt{adv}}$, we use a softmax layer on the output of the feature encoder $\mathcal{M}_{\texttt{enc}}$ and provide adversarial supervision where we classify samples to other classes to suppress identity information. Inspired by [25], we keep the "ground truth" such that identity distribution is required to be constant over all identities to train the branch. The main advantage is to remove the identification information from the Re-ID features by fooling the network with adversarial supervision. Here, we minimize the negative entropy of the predicted identity distribution and the adversarial ground truth. The loss function for the adversarial module is given by $\mathcal{L}_{\texttt{Adv-ID}}$ as follows.

$$\mathcal{L}_{\texttt{Adv-ID}} = \frac{1}{N \times K} \sum_{i=1}^{\mathcal{N}} \sum_{k=1}^{\mathcal{K}} \log \mathcal{S}(i, \theta_i)_k \tag{3}$$

where N images are fed to the network and K is the number of identities in the mini-batch $\mathcal{B}$. We set $\mathcal{S}(i, \theta_i)_k = \sigma(w_i)_k$, where σ represents the softmax layer, $\mathcal{S}(\cdot)$ repre-

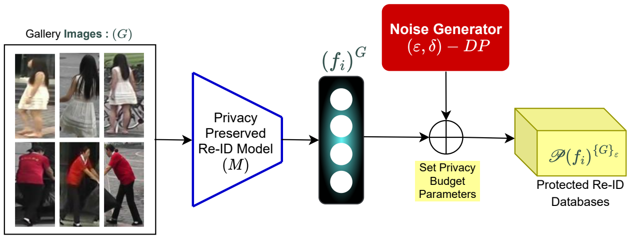
Gallery Images : (G)
Privacy Preserved Re-ID Model (M)
$(f_i)^G$
Noise Generator $(\varepsilon, \delta) - DP$
Set Privacy Budget Parameters
$\mathscr{P}(f_i)^{\{G\}_\varepsilon}$
Protected Re-ID Databases

Figure 4. Illustration of controllable privacy with differential privacy. Features $(f_i)^{\mathcal{G}}$ are extracted from a trained Re-ID model $\mathcal{M}$ for gallery $\mathcal{G}$. The image perturbation generates Gaussian noises with a given $(\epsilon, \delta - DP)$ for $(f_i)^{\mathcal{G}}$ to generate a protected gallery $\mathcal{P}(f_i)^{\mathcal{G}_\epsilon}$. A query q_i can then be used to retrieve the closest matching results from $\mathcal{P}(f_i)^{\mathcal{G}_\epsilon}$.

sents the predicted probability distribution over identities, and w_i refers to the learned weights.

4.3.3 Joint Training

We use a joint training strategy to train the network where we utilized a weighted loss function to balance the utility and privacy. Our total weighted cost function is as follows:

$$\mathcal{L}_{\texttt{total}} = \lambda_1 \mathcal{L}_{\texttt{Re-ID}} + \lambda_2 \mathcal{L}_{\texttt{De-ID}} + \lambda_3 \mathcal{L}_{\texttt{Adv-ID}} \qquad (4)$$

where λ_1, λ_2, and λ_3 are the weighted parameters whose values are empirically tuned to 1, 10^{-1} and 10^{-3}. Here, the training weights play an important role to train the network as it prevents the network to diverge from the goal of balancing utility and privacy.

4.4. Controllable Privacy via Differential Privacy

Various privacy regulations, such as GDPR [5], emphasize the need for controllable privacy in the real-world deployed system so that privacy can be adjusted as per demands. However, currently there exist no solution of controllable privacy in discriminative features. A recent comprehensive survey [27] indicates that controllable privacy is an open issue and a future research challenge. This motivates us to propose a novel controllable privacy mechanism in the Re-ID framework under a real-world setting.

As mentioned earlier, person Re-ID is a verification task (*i.e*. to determine if two images belong to the same identity) that deals with non-overlapping identities at the training phase and real-world deployment. We proposed to provide a controllable privacy to the test identities (Fig. 4). In practice, the test data from disjoint cameras is first divided into query $\mathcal{Q}$ and gallery $\mathcal{G}$, where samples in the queries $\mathcal{Q}$ are used to retrieve the closest matches from the gallery $\mathcal{G}$. Here, there is a need to protect the gallery such that PII of any individual present in the gallery cannot be leaked while preserving an acceptable retrieval performance.

Given a trained privacy-preserving Re-ID model $\mathcal{M}$, we first extract the features $(\boldsymbol{f}_i)^{\mathcal{G}}$ for the gallery $\mathcal{G}$. The proposed controllable privacy mechanism then applies differential privacy (DP) via a Gaussian noise generator with different levels of noise on $(\boldsymbol{f}_i)^{\mathcal{G}}$. Applying DP to the image domain is complex. Here, we are inspired by approximate DP where a relaxation of DP, *i.e.* (ϵ, δ)-DP, can be achieved by adding Gaussian noise. In the Gaussian mechanism, δ is greater than 0, whereas the Laplacian mechanism is known as pure DP where δ=0. In our experience, varying δ from 0.001 to 0.1 works well to achieve a good privacy-utility tradeoff. Hence, we use the Gaussian mechanism. The key idea is to add privacy-calibrated Gaussian noise that can be controlled via different privacy budget parameters.

DP Mechanism: Consider $(\boldsymbol{f}_i)^{\mathcal{G}}$ to be the features of the gallery database of the trained Re-ID model $\mathcal{M}$. A perturbation mechanism $\mathcal{P}$: $\texttt{Domain} \rightarrow \texttt{Range}$ satisfies (ϵ, δ)-differential privacy, if for any neighboring gallery database $(\boldsymbol{f}_i)^{\mathcal{G}}$, $(\boldsymbol{f}_i)^{\mathcal{G}'}$ differing on one element and for any subset of outputs $\mathcal{O} \subseteq \texttt{Range}$ holds:

$$Pr\big[\mathcal{P}\big((f_i)^{\mathcal{G}}\big) \in \mathcal{O}\,\big] \leq e^{\epsilon} Pr\big[\mathcal{P}\big((f_i)^{\mathcal{G}'}\big) \in \mathcal{O}\big] + \delta \qquad (5)$$

Here, ϵ and δ are the privacy parameters. The trade-off between accuracy and privacy can be controlled by adjusting the privacy budget parameter ϵ. A smaller privacy budget leads to more added noise which can give better privacy and hence – less privacy leakage. In our case, we empirically set $0 < \epsilon \leq 1$ and $0.001 \leq \delta \leq 0.1$.

To add privacy-calibrated noise, we employ the Gaussian mechanism by adding Gaussian noise η to the gallery features $(\boldsymbol{f}_i)^{\mathcal{G}}$. It helps protect privacy by introducing randomness with a Gaussian distribution [24]. The amount of noise required to ensure the mechanism satisfies a given privacy guarantee typically depends on how sensitive the function S where $S\big((\boldsymbol{f}_i)^{\mathcal{G}}\big) = (\boldsymbol{f}_i)^{\mathcal{G}} + \eta$; is to changes in the input and the specific distribution. The Gaussian mechanism gives a way to calibrate a zero-mean Gaussian perturbation $\mathcal{Z} \sim \mathcal{N}(0, \sigma^2 I)$ to the global L2 sensitivity given by

$$\Delta = \max_{(\boldsymbol{f}_i)^{\mathcal{G}}, (\boldsymbol{f}_i)^{\mathcal{G}'}} \Big|\Big| S\big((\boldsymbol{f}_i)^{\mathcal{G}}\big) - S(\boldsymbol{f}_i)^{\mathcal{G}'} \Big|\Big|_2 \qquad (6)$$

Sensitivity calibrates the amount of noise for gallery database. The Gaussian mechanism is (ϵ, δ)-differential privacy, if $\sigma \geq \Delta \sqrt{2log(1.25/\delta)}/\epsilon$ for any $(\epsilon, \delta) \in (0, 1)$. In our case, we use zero mean and $(\sigma_{\texttt{dataset}} - 5) \leq \sigma \leq (\sigma_{\texttt{dataset}} + 5)$ to add perturbation such that it can satisfy this relation. Here, $\sigma_{\texttt{dataset}}$ is the standard deviation of the database features.

5. Experiments

In this section, we first delineate the selected datasets and implementation details, followed by the ablation study, comparison with existing works, and qualitative analysis.

Table 1. Ablation Study on Market-1501 [44] and CUHK03-NP [20]. Here, the Utility is Re-ID accuracy and Privacy refer to identification task. mAP, Top-1 (%) and ID (%) are reported. ✓indicates the corresponding loss function is applied and otherwise ×. '–' means that the parameter is not used. Different De-ID operations is used where ⋆ represents blur, ⋆⋆ represents obfuscation, and ⋆⋆⋆ represents pixelation.

Variants	**Settings**				**Market-1501**			**CUHK03-NP**		
					Utility		Privacy	Utility		Privacy
	$\mathcal{L}_{Re-ID}$	$\mathcal{L}_{CDec}$	$\mathcal{L}_{De-ID}$	$\mathcal{L}_{Adv-ID}$	mAP	Top-1 (%)	ID (%)	mAP	Top-1 (%)	ID (%)
Baseline [ResNet-50]	✓	×	×	×	69.22	85.54	72.74	75.74	97.29	90.14
Baseline + Clean Decoder	✓	✓	×	×	69.79	86.02	74.44	76.37	97.21	90.00
Baseline + De-ID w blurring	✓	×	✓⋆	×	58.28	77.17	47.68	70.37	95.21	79.86
Baseline + De-ID w obfuscation	✓	×	✓⋆⋆	×	61.53	80.79	53.15	62.44	82.44	60.44
Baseline + De-ID w pixelation	✓	×	✓⋆⋆⋆	×	68.34	84.03	50.27	69.23	93.50	69.71
Baseline + Adv-ID	✓	×	×	✓	55.90	74.97	37.26	58.45	83.50	60.20
Proposed (1 Stage)	✓	×	✓⋆⋆⋆	✓	57.65	77.55	33.55	69.40	90.36	64.07

5.1. Datasets

Market1501 [44] is a large-scale public benchmark dataset for person Re-ID. It contains 32,668 image bounding boxes of 1,501 pedestrians captured by six different cameras including five high-resolution and one low-resolution camera. Each identity in the training set has 17.2 photos on average. We randomly divided the dataset such that 50% of the unique identities are for training and the remaining for test.

CUHK03-NP [20] contains 14,097 images of 1,467 identities. NP stands for a new protocol where it has a larger gallery with 5,332 images of 700 identities. CUHK03-NP has a smaller training set (767 identities) while the original protocol has 1467 identities.

5.2. Implementation and Evaluation Protocol

Training Details: The models are trained with NVIDIA GeForce RTX 2080 Ti GPU using TensorFlow. We use ResNet-50 [13] as the backbone with Adam Optimizer. The input images are resized to 256×128. We apply random horizontal flipping and cropping for data augmentation. The mini-batch size is set to 128 containing 32 persons with 4 images each. The initial learning rate is 0.0003 and is reduced by following an exponentially decaying training schedule until convergence is achieved.

Evaluation protocols: We use Top-1 accuracy, mean average precision (mAP) to evaluate the utility (*i.e.* Re-ID task) and classification accuracy (ID) to evaluate the privacy (*i.e.* identification task). The Re-ID for test set is calculated by matching queries against the gallery, and ID is calculated by predicting the identity of queries against the gallery. To evaluate the privacy leakage, we study whether the learned features can discriminate person's identity. Since the training and test set are disjoint, we train an identity-based softmax classifier for the learned features on a held-out set in the test set. Then, the trained classifier is used to identify the remaining test set. Here, lower ID accuracy means higher privacy preservation and vice versa. Note that higher Re-ID and lower ID accuracy lead to better performance.

5.3. Ablation Study

We conduct an ablation study to understand the effectiveness of various components in our proposed framework. The results are reported in Table 1, Table 2 and Fig. 5.

Baseline: We first experiment with widely used ResNet-50 [13] in the Re-ID domain as the baseline Re-ID encoder. The objective function is $\arg\min \mathcal{L}_{\texttt{Re-ID}}$. We can see the Top-1 Re-ID performance of 85.54% at the cost of revealing identity information of 72.74%, which can threaten privacy.

Clean Decoder: We augment the baseline network with the clean decoder where our target output is the original image. The loss function for the clean decoder is given by $\mathcal{L}_{\texttt{CDec}}$ where $\mathcal{L}_{\texttt{CDec}} = \sum_{i=1}^{N} \|\boldsymbol{x}_i - \hat{\boldsymbol{x}}_i\|_2^2$. Here, $\boldsymbol{x}_i$ is the original input image, and the reconstructed images are denoted by $\hat{\boldsymbol{x}}_i$, and N is the number of images. When we add a clean decoder to reconstruct the original input image, we see an improvement of +1.7% in ID accuracy for Market-1501. This shows that the network is able to exploit deep-appearance features to gain performance.

De-identified (De-ID) Decoder: We use different de-identification mechanisms, such as pixelation, obfuscation, and blurring, to inject the de-identified information and provide respective de-identification images as supervision to the decoder. The objective function of the de-identified decoder is $\arg\min \mathcal{L}_{\texttt{De-ID}}$ given in Eqn (2). Here, we optimize the baseline encoder with a deidentified decoder using $\arg\min \mathcal{L}_{\texttt{Re-ID}}$ and $\arg\min \mathcal{L}_{\texttt{De-ID}}$. We observe that the pixelation operation achieves a better trade-off between privacy and utility as compared to other De-ID operations. Specifically, the ID accuracy reduced by -22.47% whereas the Re-ID performance is reduced by -0.88 mAP. In comparison, the obfuscation and blurring mechanism greatly reduce the Re-ID performance.

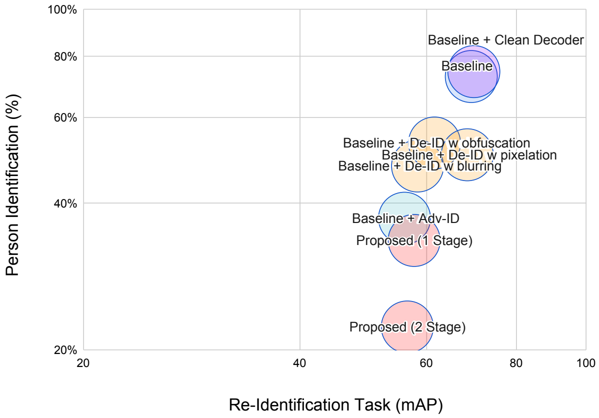

Figure 5. Comparison of the proposed methods, their variants, and the baseline method on the privacy-utility trade-off.

Adversarial-Identity (Adv-ID) Module: Here, we utilize the softmax layer and provide adversarial supervision through a constant identity distribution as the ground truth. The objective function is set to $\arg\min \mathcal{L}_{\texttt{Adv-ID}}$ given in Eqn (3). Here, we find that this component helps to remove identifiable information (*i.e.* -37.18% in Market-1501) and preserves privacy. However, it is adversely impacting the Re-ID accuracy. Here, we use a weighting mechanism in the full proposed approach to get the advantage from this component effectively.

Privacy-Preserved Re-ID Model (Proposed (1 Stage)): The proposed model combines a Re-ID encoder, a de-identified decoder (through pixelation), and an adversarial module. We use a joint training strategy and the overall objective function is $\arg\min \mathcal{L}_{\texttt{total}}$ (c.f. Eqn (4)), where we give higher weight to Re-ID loss and less weight to other losses (*i.e.* $\lambda_2 = 0.1$ and $\lambda_3 = 0.001$) to bring the balance between utility and privacy.

Privacy-Preserved Re-ID Model with Controllable Privacy (Proposed (2 Stage)): We utilize the power of DP to control privacy through different privacy budget parameters. Different values of the privacy budget parameters (*i.e.* ϵ = 0.1, 0.5, and 1) are selected to demonstrate this. Our preliminary study shows that the δ is more stable within the range of 0.01 to 0.001. We compare the performance of both baseline and our proposed privacy preserved Re-ID model (*i.e.* Proposed (1 Stage)) when the controllable privacy mechanism is applied. As shown in Table 2, we can observe that adding differential privacy-based perturbation to the proposed model can prevent the leakage of identity information more instead of directly adding it to the baseline models. For example, at ϵ=0.5, person identification performance for baseline is dropped by only -12.34%, whereas on our method the performance dropped is -50.44% and the Re-ID accuracy is still comparable. Comparing Proposed (1 stage) with Proposed (2 stage), we observed good privacy preservation (33.55% vs 22.3% in accuracy) without heavily compromised on the utility (57.65 vs 56.4 in mAP).

Table 2. Performance of baseline and proposed method with controllable privacy. mAP, Top-1 (%), and ID (%) results are reported on Market-1501 dataset. The performance loss against baseline are shown in BLUE).

Methods	mAP	Top-1 (%)	ID (%)
Baseline	69.22	85.54	72.74
Baseline + DP ($\epsilon = 0.1$)	55.0 (-14.22)	71.9 (-13.64%)	49.0 (-23.74%)
Baseline + DP ($\epsilon = 0.5$)	62.7 (-6.52)	79.5 (-6.04%)	60.4 (-12.34%)
Baseline + DP ($\epsilon = 1.0$)	63.5 (-5.72)	80.9 (-4.64%)	71.1 (-1.64%)
Proposed (1 Stage)	57.65 (-11.57)	77.55 (-7.99%)	33.55 (-39.19%)
Proposed (2 Stage, $\epsilon = 0.1$)	51.5 (-17.72)	69.8 (-15.74%)	14.1 (-58.64%)
Proposed (2 Stage, $\epsilon = 0.5$)	56.4 (-12.82)	75.3 (-10.24%)	22.3 (-50.44%)
Proposed (2 Stage, $\epsilon = 1.0$)	56.9 (-12.32)	75.8 (-9.74%)	29.9 (-42.84%)

5.4. Comparison with SOTA methods

We compare our proposed method with Re-ID baselines with various backbones and privacy-preserving based SOTA methods on Market-1501 (Table 3) and CUHK03-NP (Table 4). All SOTA privacy-preserving methods focus only on preserving the Re-ID performance and not reporting the privacy leakage information.

First, we demonstrate privacy leakage in existing backbone networks (*i.e.* MobileNet [16], ResNet-50, and ResNet-101 [13]) and compare the identification accuracy with the proposed method. Results shows that 72.74% (for Market-1501) and 90.14% (for CUHK03-NP) of identity information gets leaked with ResNet-50 backbone, whereas the proposed (1-stage) achieved better privacy protection (*i.e.* 33.55% for market-1501 and 64.07% for CUHK03-NP). The protection can be further enhanced with controllable privacy mechanism (*i.e.* Proposed (2-stage)) where the identification accuracy is 23.30% and 41.14% for Market-1501 and CUHK03-NP, respectively, when $\epsilon = 0.5$. This clearly demonstrates that our learned features have less PII to identify individuals as compared to baseline.

Compared against privacy-preserving SOTA models, including DP-SGD [1], federated learning [45, 46], Face-blur MuDeep [9], and PIS [10], our proposed method outperforms all of the compared models while achieving a better trade-off between Re-ID performance and privacy. For example, Proposed (1 Stage) achieved 69.4 mAP in re-identification task on CUHK03-NP while best performing SOTA only achived 50.4 mAP. Even with further privacy protection with differential privacy (*i.e.* Proposed (2 Stage) with $\epsilon = 0.5$), we still achieved 62.45 mAP. Note that the privacy improvement at this stage is almost at 50% improvement. The experiments empirically demonstrates the effectiveness of the proposed model.

Table 3. Comparisons with different baselines and privacy-based SOTA methods on the Market-1501 dataset. mAP, Top-1 (%) and Identification accuracy (%) are reported. – represents privacy ID is not reported by these methods.

Methods	**Models**	**Backbone**	**Utility**		**Privacy**
			mAP	Top-1 (%)	ID (%)
Without Privacy Preserving	Mobile-Net [16]	MobileNet-V1	55.52	75.18	66.15
	ResNet-101 [13]	ResNet-101	64.68	83.22	68.32
	ResNet-50 [13]	ResNet-50	69.22	86.02	72.74
Privacy Preserving	DP-SGD [1]	ResNet-50	4.5	17.6	–
	Federated-by-Camera [46]	ResNet-50	36.57	61.13	–
	Face-blur MuDeep [9]	4L ConvNet	44.8	69.6	–
	PIS [10]	ResNet-50	51.9	74.9	–
	Proposed (1 Stage)	ResNet-50	57.65	77.55	33.55
Controllable Privacy	**Proposed (2 Stage,** $\epsilon = 1.0$**)**	ResNet-50	56.90	75.80	29.90
	Proposed (2 Stage, $\epsilon = 0.5$**)**	ResNet-50	56.40	75.30	22.30
	Proposed (2 Stage, $\epsilon = 0.1$**)**	ResNet-50	51.50	69.80	14.10

Table 4. Comparisons with different baselines and privacy-based SOTA methods on the CUHK03-NP dataset. mAP, Top-1 (%), and Identification accuracy (%) are reported.

Methods	**Models**	**Backbone**	**Utility**		**Privacy**
			mAP	Top-1 (%)	ID (%)
Without Privacy Preserving	Mobile-Net [16]	MobileNet-V1	64.29	93.36	84.50
	ResNet-101 [13]	ResNet-101	66.27	93.29	87.21
	ResNet-50 [13]	ResNet-50	75.74	97.29	90.14
Privacy Preserving	Federated-by-Camera [46]	ResNet-50	11.11	11.21	–
	Federated-by-Identity [45]	ResNet-50	47.39	51.71	–
	Face-blur MuDeep [9]	4L ConvNet	22.2	23.3	–
	Face-blur HACNN [9]	Inception	32.9	32.4	–
	Face-blur PCB [9]	ResNet-50	50.4	51.1	–
	Proposed (1 Stage)	ResNet-50	69.40	90.36	64.07
Controllable Privacy	**Proposed (2 Stage,** $\epsilon = 1.0$**)**	ResNet-50	67.12	87.23	62.47
	Proposed (2 Stage, $\epsilon = 0.5$**)**	ResNet-50	62.45	79.32	41.14
	Proposed (2 Stage, $\epsilon = 0.1$**)**	ResNet-50	55.21	71.75	28.79

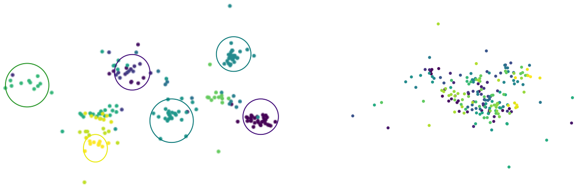

(a) Baseline (b) Proposed (Ours)

Figure 6. Visualization of t-SNE plots on Market-1501. Each color indicates a unique identity.

5.5. Qualitative Comparison

Feature Visualization with t-SNE. We randomly select a set of identities from the test set and visualize the corresponding features using t-SNE [26] in Fig. 6. The baseline features (*i.e*. Fig. 6a) show the same identity are strongly clustered due to identity-dependent information. Hence, it is subjected to identity information breach of the observed individual. In contrast, the features from proposed method (*i.e*. Fig. 6b) are more uniformly spaced, revealing that they are more identity invariant and contain less personal information. This visualization shows the efficacy of the proposed proposed model over the baseline features, and it is consistent with the quantitative results from Section 5.4.

Reconstructed Images. The first row in Fig. 7 shows the original images and the second row shows the reconstructed images with the help of a clean decoder (where original images are given as supervision). Here, we can clearly see that an adversary can decode PII, such as clothing, gender, body type, *etc*., from Re-ID features. This shows that identity related features are somehow encoded. In the third row, we show results with a de-identified decoder which demonstrates that appearance information to identify individuals is suppressed. Furthermore, in the fourth row, we can see that identifiable information cannot be decoded from our proposed model. This validates that our proposed model has a strong capability to suppress identity-related information and to learn the privacy-preserved Re-ID features.

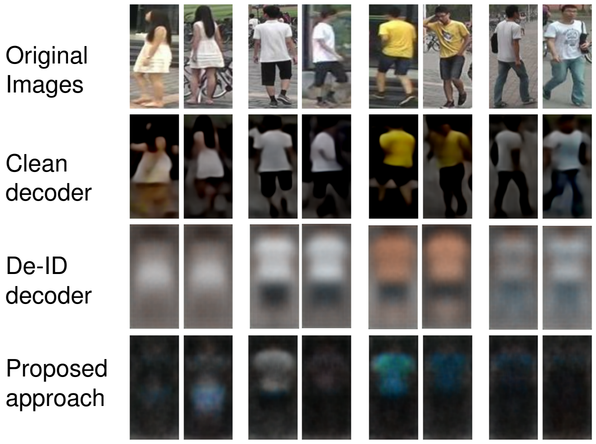

Figure 7. Samples of the reconstructed images from Market-1501.

6. Conclusion

This work proposes a controllable privacy-preserving model to learn robust privacy-preserved Re-ID features. We utilized a De-ID decoder and adversarial supervision module to suppress the identity information during the model training stage. To achieve controllable privacy, we apply the DP mechanism on the feature space to control the identity information based on the different privacy budgets. Our results demonstrate that the learned privacy-preserved Re-ID features have a strong capability to balance the tradeoff between utility and privacy. For future work, one critical direction is to improve the utility preservation when a model tries to suppress the encoded PII. Another direction is to study the feasibility of incorporating perturbed images via the DP mechanism during Re-ID model training.

Acknowledgement

This research is supported by the National Research Foundation, Singapore under its Strategic Capability Research Centres Funding Initiative. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

References

[1] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *CCS*, pages 308–318, 2016.

[2] Shafiq Ahmad, Gianluca Scarpellini, Pietro Morerio, and Alessio Del Bue. Event-driven re-id: A new benchmark and method towards privacy-preserving person re-identification. In *WACV*, pages 459–468, 2022.

[3] Abdullah Alshaibani and Alexander J Quinn. Pterodactyl: Two-step redaction of images for robust face deidentification. In *AAAI*, volume 9, pages 27–34, 2021.

[4] Igor Barros Barbosa, Marco Cristani, Alessio Del Bue, Loris Bazzani, and Vittorio Murino. Re-identification with RGB-D sensors. In *ECCV Workshops*, volume 7583 of *Lecture Notes in Computer Science*, pages 433–442. Springer, 2012.

[5] Eduard Barnoviciu, Veta Ghenescu, Serban-Vasile Carata, Marian Ghenescu, Roxana Mihaescu, and Mihai Chindea. Gdpr compliance in video surveillance and video processing application. In *SpeD*, pages 1–6, 2019.

[6] Mahawaga Arachchige Pathum Chamikara, Peter Bertok, Ibrahim Khalil, Dongxi Liu, and Seyit Camtepe. Privacy preserving face recognition utilizing differential privacy. *Computers & Security*, 97:101951, 2020.

[7] Hang Cheng, Huaxiong Wang, Ximeng Liu, Yan Fang, Meiqing Wang, and Xiaojun Zhang. Person re-identification over encrypted outsourced surveillance videos. *IEEE Transactions on Dependable and Secure Computing*, 18(3):1456–1473, 2019.

[8] Warren B Chik. The singapore personal data protection act and an assessment of future trends in data privacy reform. *Computer Law & Security Review*, 29(5):554–575, 2013.

[9] Julia Dietlmeier, Joseph Antony, Kevin McGuinness, and Noel E O'Connor. How important are faces for person re-identification? In *ICPR*, pages 6912–6919, 2021.

[10] Shuguang Dou, Xinyang Jiang, Qingsong Zhao, Dongsheng Li, and Cairong Zhao. Towards privacy-preserving person re-identification via person identify shift. *arXiv preprint arXiv:2207.07311*, 2022.

[11] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *EUROCRYPT*, volume 4004 of *Lecture Notes in Computer Science*, pages 486–503. Springer, 2006.

[12] Jianyuan Guo, Yuhui Yuan, Lang Huang, Chao Zhang, Jin-Ge Yao, and Kai Han. Beyond human parts: Dual part-aligned representations for person re-identification. In *ICCV*, pages 3642–3651, 2019.

[13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.

[14] Muqing He, Deqing Zou, Weizhong Qiang, Wenbo Wu, Shouhuai Xu, Xianjun Deng, and Hai Jin. Utility-prioritized differential privacy for quantitative biomedical data. *Journal of Circuits, Systems and Computers*, 31(13):2250236:1–2250236:23, 2022.

[15] Steven Hill, Zhimin Zhou, Lawrence K Saul, and Hovav Shacham. On the (in) effectiveness of mosaicing and blurring as tools for document redaction. *Proc. Priv. Enhancing Technol.*, 2016(4):403–417, 2016.

[16] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. MobileNets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.

[17] Kajal Kansal, AV Subramanyam, Zheng Wang, and Shinichi Satoh. Hierarchical attention image-text alignment network for person re-identification. In *2021 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*, pages 1–6. IEEE, 2021.

[18] Kajal Kansal, A Venkata Subramanyam, Zheng Wang, and Shin'ichi Satoh. Sdl: Spectrum-disentangled representation learning for visible-infrared person re-identification. *IEEE TCSVT*, 30(10):3422–3432, 2020.

[19] Tao Li and Lei Lin. AnonymousNet: Natural Face De-Identification With Measurable Privacy. In *CVPR Workshops*, pages 56–65, 2019.

[20] Wei Li, Rui Zhao, Tong Xiao, and Xiaogang Wang. DeepReID: Deep filter pairing neural network for person re-identification. In *CVPR*, pages 152–159, 2014.

[21] Wei Li, Xiatian Zhu, and Shaogang Gong. Harmonious attention network for person re-identification. In *CVPR*, pages 2285–2294, 2018.

[22] Ji Lin, Liangliang Ren, Jiwen Lu, Jianjiang Feng, and Jie Zhou. Consistent-aware deep learning for person re-identification in a camera network. In *CVPR*, pages 5771–5780, 2017.

[23] Yutian Lin, Liang Zheng, Zhedong Zheng, Yu Wu, Zhilan Hu, Chenggang Yan, and Yi Yang. Improving person re-identification by attribute and identity learning. *Pattern Recognition*, 95:151–161, 2019.

[24] Fang Liu. Generalized gaussian mechanism for differential privacy. *IEEE Transactions on Knowledge and Data Engineering*, 31(4):747–756, 2018.

[25] Yu Liu, Fangyin Wei, Jing Shao, Lu Sheng, Junjie Yan, and Xiaogang Wang. Exploring disentangled feature representation beyond face identification. In *CVPR*, pages 2080–2089, 2018.

[26] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008.

[27] Blaž Meden, Peter Rot, Philipp Terhörst, Naser Damer, Arjan Kuijper, Walter J Scheirer, Arun Ross, Peter Peer, and Vitomir Štruc. Privacy–enhancing face biometrics: A comprehensive survey. *IEEE Transactions on Information Forensics and Security*, 16:4147–4183, 2021.

[28] Shunmei Meng, Shaoyu Fan, Qianmu Li, Xinna Wang, Jing Zhang, Xiaolong Xu, Lianyong Qi, and Md Zakirul Alam Bhuiyan. Privacy-aware factorization-based hybrid recommendation method for healthcare services. *IEEE Transactions on Industrial Informatics*, 18(8):5637–5647, 2022.

[29] Seyed Ali Miraftabzadeh, Paul Rad, Kim-Kwang Raymond Choo, and Mo Jamshidi. A privacy-aware architecture at

the edge for autonomous real-time identity reidentification in crowds. *IEEE Internet of Things Journal*, 5(4):2936–2946, 2017.

[30] Marina Paolanti, Luca Romeo, Daniele Liciotti, Rocco Pietrini, Annalisa Cenci, Emanuele Frontoni, and Primo Zingaretti. Person re-identification with RGB-D camera in top-view configuration through multiple nearest neighbor classifiers and neighborhood component features selection. *Sensors*, 18(10):3471, 2018.

[31] Yuqing Peng, Wei Li, Yingjun Li, Yixin Pei, and Yongfang Guo. Multi-task person re-identification via attribute and part-based learning. *Multimedia Tools and Applications*, 81(8):11221–11237, 2022.

[32] Nan Pu, Zhun Zhong, Nicu Sebe, and Michael S Lew. A memorizing and generalizing framework for lifelong person re-identification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.

[33] Luc Rocher, Julien M Hendrickx, and Yves-Alexandre De Montjoye. Estimating the success of re-identifications in incomplete datasets using generative models. *Nature communications*, 10(1):1–9, 2019.

[34] Chi Su, Jianing Li, Shiliang Zhang, Junliang Xing, Wen Gao, and Qi Tian. Pose-driven deep convolutional model for person re-identification. In *ICCV*, pages 3960–3969, 2017.

[35] Yifan Sun, Liang Zheng, Yali Li, Yi Yang, Qi Tian, and Shengjin Wang. Learning part-based convolutional features for person re-identification. *IEEE TPAMI*, 43(3):902–917, 2019.

[36] Jhanvi Uday and Mohona Ghosh. Safeguarding GeoLocation for Social Media with Local Differential Privacy and L-Diversity. In *Security, Privacy and Data Analytics*, pages 17–31. Springer, 2022.

[37] Haochen Wang, Jiayi Shen, Yongtuo Liu, Yan Gao, and Efstratios Gavves. Nformer: Robust person re-identification with neighbor transformer. In *CVPR*, pages 7297–7307, 2022.

[38] Guile Wu and Shaogang Gong. Decentralised learning from independent multi-domain labels for person re-identification. In *AAAI*, pages 2898–2906, 2021.

[39] Suncheng Xiang, Hao Chen, Wei Ran, Zefang Yu, Ting Liu, Dahong Qian, and Yuzhuo Fu. Deep multimodal representation learning for generalizable person re-identification. *Machine Learning*, pages 1–19, 2023.

[40] Zhengwei Yang, Xian Zhong, Zhun Zhong, Hong Liu, Zheng Wang, and Shin'ichi Satoh. Win-win by competition: Auxiliary-free cloth-changing person re-identification. *IEEE Transactions on Image Processing*, 2023.

[41] Mang Ye, Chao Liang, Yi Yu, Zheng Wang, Qingming Leng, Chunxia Xiao, Jun Chen, and Ruimin Hu. Person reidentification via ranking aggregation of similarity pulling and dissimilarity pushing. *IEEE Transactions on Multimedia*, 18(12):2553–2566, 2016.

[42] Mang Ye, Jianbing Shen, Gaojie Lin, Tao Xiang, Ling Shao, and Steven CH Hoi. Deep learning for person re-identification: A survey and outlook. *IEEE TPAMI*, 44(6):2872–2893, 2021.

[43] Bowen Zhao, Yingjiu Li, Ximeng Liu, Hwee Hua Pang, and Robert H Deng. FREED: An efficient privacy-preserving solution for person re-identification. In *DSC*, pages 1–8, 2022.

[44] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *ICCV*, pages 1116–1124, 2015.

[45] Weiming Zhuang, Xin Gan, Yonggang Wen, and Shuai Zhang. Optimizing performance of federated person re-identification: Benchmarking and analysis. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 2022.

[46] Weiming Zhuang, Yonggang Wen, Xuesen Zhang, Xin Gan, Daiying Yin, Dongzhan Zhou, Shuai Zhang, and Shuai Yi. Performance optimization of federated person re-identification via benchmark analysis. In *ACM MM*, pages 955–963, 2020.