

# Holistic Guidance for Occluded Person Re-Identification

Madhu Kiran<sup>1</sup>

madhu\_sajc@hotmail.com

R Gnana Praveen<sup>1</sup>

praveengnan.24@gmail.com

Le Thanh Nguyen-Meidine<sup>1</sup>

le-thanh.nguyen-meidine.1@ens.etsmtl.ca

Soufiane Belharbi<sup>1</sup>

soufiane.belharbi.1@ens.etsmtl.ca

Louis-Antoine Blais-Morin<sup>2</sup>

lablaimorin@GENETEC.COM

Eric Granger<sup>1</sup>

Eric.Granger@etsmtl.ca

<sup>1</sup> LIVIA, Dept. of Systems Engineering

École de technologie supérieure  
Montreal, Canada

<sup>2</sup> Genetec Inc.

Montreal, Canada

---

## Abstract

In real-world video surveillance applications, person re-identification (ReID) suffers from the effects of occlusions and detection errors. Despite recent advances, occlusions continue to corrupt the features extracted by state-of-art CNN backbones and thereby deteriorate the accuracy of ReID systems. To address this issue, methods in the literature rely on an additional costly process, such as pose estimation, where pose maps provide supervision to focus on visible parts of occluded regions. In contrast, we introduce a Holistic Guidance (HG) method that relies on holistic (or non-occluded) data and its distribution in the dissimilarity space to train the CNN backbone on an occluded dataset. This method is motivated by our empirical study, where the distribution of pairwise between-class and within-class matching distances (Distribution of Class Distances or DCDs) between images has considerable overlap in occluded datasets compared to holistic datasets. Hence, our HG method employs this discrepancy in DCDs of both datasets for joint learning of a student-teacher model to produce an attention map that focuses primarily on visible regions of the occluded images. In particular, features extracted from both datasets are jointly learned using the student model to produce an attention map that allows dissociating visible regions from occluded ones. Additionally, a joint generative-discriminative CNN backbone is trained using a denoising autoencoder such that the system can self-recover from occlusions. Extensive experiments on several challenging public datasets indicate that the proposed approach can outperform state-of-the-art methods on both occluded and holistic datasets. **Our code is available<sup>1</sup>.**

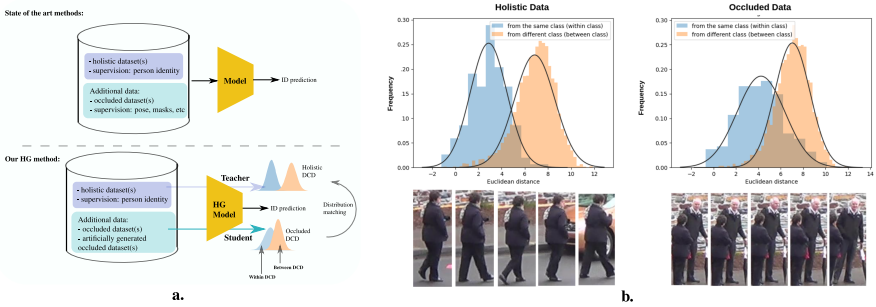


Figure 1: **(a)** An illustration of approaches to address occlusion in person ReID during training. **Top:** State-of-the-art models require additional supervision and occluded datasets. **Bottom:** Our proposed HG method requires no additional supervision but relies only on an additional holistic dataset for reference to non-corrupted features. **(b)** Examples of Class Distance Distributions of Duke-MTMC (left) and Occluded-Duke-MTMC (right) datasets measured in the distance space. The **blue DCD** shows the within-class distribution, while the **orange DCD** shows the between-class distribution. For Occluded-Duke-MTMC, within-class distances are relatively high and overlap with between-class distances.

## 1 Introduction

Person Re-Identification (ReID) systems seek to associate individuals captured over a distributed set of non-overlapping camera viewpoints. This key visual recognition task has recently drawn significant attention due to its wide range of applications, e.g., autonomous driving, pedestrian tracking, sports analytics, and video surveillance [11, 28, 29, 40]. Despite the recent progress with deep learning (DL), person ReID remains a challenging task in real-world applications due to the non-rigid structure of the human body, variability of capture conditions (e.g., illumination, scale, motion blur), in addition to person detection issues like miss-alignment, background clutter, and occlusion [16, 26, 52].

This paper focuses on the occlusion issue for person ReID, a challenge that has attracted much attention [12, 15, 16, 26, 53, 56]. When bounding boxes or regions of interest (ROI) are occluded, the CNN backbone extracts noisy features, leading to pairwise matching errors between query and reference ROIs and poor ReID accuracy for the occluded class. Since occlusions are diverse in color, shape, and size, extracting features from the entire ROI can potentially corrupt the global representation.

Several authors have attempted to address occluded person ReID by using pedestrian detectors that can additionally refine person ROIs [41]. Other methods follow an intuitive solution of masking occluded regions, extracting occlusion aware features, or applying weights and masks to occluded regions, applying body masks based on pose estimation, etc. [4, 17, 26]. With these state-of-the-art methods, the external mechanism for mask generation adds a considerable time complexity during run-time. Unlike these methods, we propose a new method that only requires person identity labels as supervision and does not rely on additional supervision such as pose estimation. This provides robustness to occlusion in person ReID, yet lower complexity during inference.

Our method is motivated by the fact that the distribution of features learned from occluded and holistic<sup>2</sup> ReID datasets are different. The dissimilarity space can be regarded as

<sup>2</sup>In the person ReID, the term "holistic" refers to image data that contains the full body of a person in both query

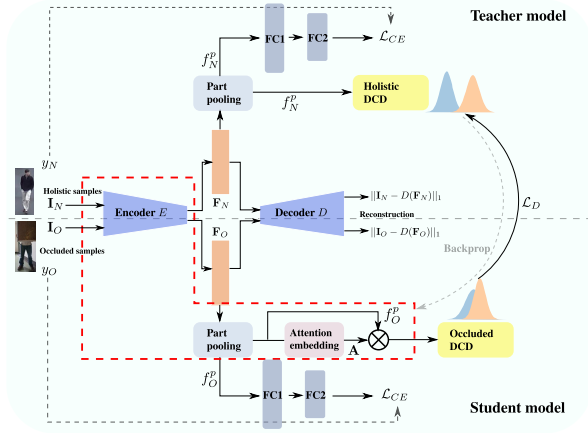


Figure 2: Our proposed HG method where a teacher model uses a holistic data distance distribution to train the student network (trained on artificially occluded or real occluded samples) such that it can accurately recognize persons appearing in occluded images.

space defined by dissimilarity coordinates, and CNN features are transformed into that space by computing pairwise matching distances for within-class and between-class samples in a given batch. It has been shown to successfully learn to separate feature representations for data that is noisy and overlapping [8, 18]. Occluded person ReID is a good example of a problem with class overlap. Inspired by [4, 9, 10, 18], we consider the dissimilarity space to capture the discrepancy between images in occluded and holistic datasets.

Fig. 1 shows the distribution of within-class and between-class distances (DCDs) among pairs of samples extracted from occluded and holistic datasets. We note two aspects: (1) **Within-class DCDs**: In an occluded dataset, this distance tends to be greater, with high variance, compared to holistic or non-occluded cases. Such distances are normally expected to be lower due to the similarity among samples within the same class. (2) **Overlap within- and between-class DCDs**: Large within-class DCDs are caused by samples being pushed away from the same class, allowing for substantial overlap with samples from other classes. The overlap between samples of different classes is more likely to impede discrimination among classes, leading to poor recognition. Note that both datasets used to generate Fig. 1 – Duke-MTMC (Holistic) [29] and Occluded-Duke [26] – are from the same domain, but occlusions still cause such discrepancies. Based on these observations of DCDs, and considering the poor performance of models over occluded datasets, we hypothesize that occlusion is a potential source of corruption for feature representations in person ReID tasks.

This paper proposes a Holistic Guidance (HG) student-teacher network that relies on the distribution of holistic data in the dissimilarity space to train a student (CNN backbone) on an occluded dataset. The discrepancy of within- and between-class DCDs across datasets allows the network to extract features on occluded samples while simultaneously maintaining a good between- and within-class separation. Models trained on occluded data tend to overfit due to class overlap, so we advocate for using guidance from non-corrupted features of a larger holistic data in the dissimilarity space to mitigate this issue. Although both datasets can have different identities, transforming the samples to the dissimilarity space translates into a binary classification problem [4, 18]. In practice, such learning scenarios could be

achieved by using a single holistic dataset, and by building an artificially (augmented) occluded dataset. A second alternative involves using a training set consisting of real occluded samples. Our method performs HG since it relies on features with properties learned from holistic data to guide the CNN backbone in learning features of an occluded dataset.

Our method (shown in Fig.2) is also comprised of a shared generative model (i.e., a denoising auto-encoder) that is trained simultaneously on both datasets to enable self-recovery from occlusion. The student model has an additional embedding for producing an attention map, allowing the partial or local features to re-focus attention only on visible body-part features while ignoring the occluded regions that cause distribution discrepancy. Several authors have proposed generative models for person ReID [58, 48], mainly for GAN based data augmentation. In contrast, we introduce a denoising auto-encoder as the CNN backbone for our HG student-teacher network, allowing to self-recover in cases of occlusion.

**Main contributions:** (1) A novel HG student-teacher network that relies on the distribution of holistic data in the dissimilarity space to train a CNN backbone on the occluded dataset. (2) To motivate our HG method, we show that within-class DCDs of Occluded-ReID datasets overlap with between-class DCDs by a larger margin than holistic ReID datasets, even in cases where the occluded dataset is a subset of a holistic one. (3) Extensive experiments were performed on challenging Occluded [26, 51], Partial [26, 47], and Holistic [29, 44] ReID datasets show that our HG method can outperform many SOTA methods.

## 2 Related Work in Person-ReID

**Image-Based Methods:** Siamese Networks have first been used in [39] that employs three Siamese sub-networks for deep feature learning. Most of the further work based on deep-architecture ReID [11, 9, 8, 25, 55] approaches introduce an end-to-end ReID framework, where both feature embedding and metric learning have been investigated. A few attention-based approaches for deep ReID [21, 61, 43] address misalignment challenges by incorporating a regional attention sub-network into a base re-ID model. With part-based methods, local features are extracted from different regions to enhance the discriminative power of the features. Suh et al. [52] extracted parts from the feature map and trained each part with separate classifiers. [20] used parts method to extract local features. In addition to this, other methods such as [53, 43] have used part pooling with attention for a refined partial feature.

**Occluded Person ReID:** Occluded person ReID is different from person ReID or holistic person ReID because during test time, the probe images are often occluded as in real-world applications. Hence [51] have proposed to use a binary classifier to classify the images as occluded or not to distinguish occluded ones from holistic images. [26] considered using pose guided feature alignment to align part features or local features. Similar to [26], other works align local or part features by pose estimation like [57, 45]. Similar to this, [56] use pose estimator to help predict key-points on body parts and use graph-based methods. The main disadvantage of pose-aligned methods is that they suffer from requiring additional supervision like the pose estimation step, which is error prone [42]. Recently, [12] use a method where, in the first step, they obtain discriminative features by using pose guided attention and then mine for visible parts. [16] have used foreground-background mask instead of the pose. Recently, [23] used only identity information to learn occluded ReID with part aware transformers. Note that they add additional complexity with the transformers used for attention. [19] is an other approach using only identity information, where Occluded ReID is considered as a set matching task to make invariant to occlusion of different parts.

Similar to our method, Zhuo et al. [52] employ a student-teacher model. In the first stage,

their teacher network learns from a holistic dataset, where both identification and supervised saliency is learned using a salient object detector. However, it differs from our approach in that: (1) it does not learn any supervised salient region detection, and (2) it allows recovering from small occlusions by using a denoising auto-encoder as a CNN backbone.

### 3 Proposed Approach

A Holistic Guidance (HG) method is introduced for unsupervised learning of attention maps in Occluded-ReID, without the need for any external guidance such as pose or segmentation maps. A student-teacher network is proposed, where the teacher network relies on holistic images or non-occluded images to teach the student network the DCD of holistic features. This allows the student to learn an attention map such that, when applied to occluded images, results in uncorrupted features like those of the teacher. More specifically, a joint generative and discriminative backbone is trained with a denoising autoencoder for the student-teacher model. It simultaneously learns to match image pairs while reconstructing images.

**Problem Formulation:** Let  $I_O$  and  $I_N$  denote input images from the occluded and holistic dataset, respectively.  $N$  and  $O$  denote the components from teacher and student trained from holistic data and occluded data, respectively. Let  $y_O$  and  $y_N$  denote the identity labels of occluded and holistic datasets.  $\mathbf{F}_O$  and  $\mathbf{F}_N$  represent the global feature maps of occluded and holistic datasets, which are obtained from the shared encoder  $E$ . The local or part-based features  $f_O^i, f_N^i$  are produced by applying a pooling function on  $\mathbf{F}_O$  and  $\mathbf{F}_N$ . By minimising the discrepancy between-class distance distribution of  $f_O^i, f_N^i$ , we intend to learn an attention map  $A^i$  that is applied on the  $f_O^i$  to obtain features  $f_a^i$  for  $i = 1, \dots, p$  for each part. During testing,  $\psi$ , which is a concatenation of global and local features, allows extracting features for matching from the gallery and query using a distance function, and to retrieve the identity.

**Robust Backbone Model:** We propose a joint learning framework of a denoising auto-encoder along with the classification networks to be robust to occlusions for person ReID. The input images are augmented by adding small noise using random erasing data augmentation, while the reconstruction loss is obtained using actual images. The encoder  $E$  is shared between denoising auto-encoder and ReID classification layers. In order to obtain robust features with both Generative and Discriminative properties, we reconstruct the input images using a Decoder on the embedding  $\mathbf{F}_N$  and  $\mathbf{F}_O$ . Hence  $E$  and  $D$  together form a denoising auto-encoder. We have used a denoising autoencoder in order to exploit the full potential of the generative capability to take full advantage of the class distributions of the holistic data-set. Let  $\hat{\mathbf{I}}_c$  represent the reconstructed image, and  $\mathbf{F}_c = E(\mathbf{I}_c)$  be the latent feature representation of the encoder, where  $c \in \{N, O\}$  (holistic and occluded images) and  $\mathbf{I}_c$  is the input image. The size of  $\mathbf{F}_c$  is  $B \times C \times w \times h$ , where  $B$  is the batch,  $C$  is the number of output channels of the encoder  $E$  and  $w, h$  width and height of the feature map.  $\mathbf{F}_c$  can also be referred to as the latent feature representation of the auto-encoder. A part-based pooling is applied on  $\mathbf{F}_c$  to obtain  $p$  parts of features. Our part pooling method used is similar to [B3] where a 2D feature map is split into horizontal stripes of  $p$  parts. Global average pooling is then performed on each of the  $p$  parts to obtain  $p$  feature vectors of size  $C$ . Each feature vector of  $p$  parts is assigned to a unique classifier, resulting in  $p$  classifiers trained using identity labels of the corresponding datasets.

The predicted output for each given image  $\mathbf{I}_c$  from the classifier is  $\hat{y}_{i,c}$ , where  $i = 1, \dots, p$  parts. The identity prediction loss function over all the parts is:

$$\mathcal{L}_{CE,c} = \frac{1}{K} \sum_i^K -\log \left( \frac{\exp(\mathbf{W}_{y_i}^T \mathbf{x}_i + b_{y_i})}{\sum_{j=1}^N \exp(\mathbf{W}_j^T \mathbf{x}_i + b_j)} \right) \quad (1)$$

where  $\mathcal{L}_{\text{CE},c}$  is the cross entropy loss,  $K$  is the batch size, class label  $y_i \in \{1, 2, \dots, N\}$  is associated with  $i^{\text{th}}$  training image.  $\mathbf{W}_j$  and  $b_j$  are the weights and bias of last fully connected layer for class  $y$ . Similarly  $\mathbf{W}_j$  and  $b_j$  are the weights and bias of the  $j^{\text{th}}$  class. The denoising auto-encoder is learned using reconstruction loss between the reconstructed and original images before random erasing augmentation [49], which acts like noise added in denoising auto-encoders. We propose this inspired by [6] to self recover from occlusion. Normally in [6] an autoencoder is learnt to generate non occluded face image from artificially occluded face image which is then post processed for recognition. But in our case since we use a joint representation we expect the joint feature to have the self recovering properties. The reconstruction loss of the autoencoder is given by,

$$\mathcal{L}_{\text{recon},c} = \mathbb{E} [\|\mathbf{I}_c - D(\mathbf{F}_c)\|_1], \quad (2)$$

where  $D$  denotes the decoder function of the denoising auto-encoder. We do not use the reconstructed image for any other steps. However, the reconstruction loss is optimised in order to enable the deep features to have generative properties and to self-recover from occlusion. The total loss for the joint learning of generative discriminative learning is:

$$\mathcal{L}_{\text{joint},c} = \mathcal{L}_{\text{CE},c} + \lambda \mathcal{L}_{\text{recon},c}, \quad (3)$$

where  $\lambda$  is the trade-off parameter.

**Student-Teacher Model:** A student-teacher model with the proposed backbone is shared between the teacher and student. In addition to the backbone, the student model carries an embedding to produce attention maps. From Fig. 1, it can be observed that a DCD obtained by comparing extracted deep features of occluded in-class images overlap with those of out-of-class distances by a large margin in comparison with the holistic dataset. The overlap of the DCD indicates corrupted features as in-class distance distribution must have good separation from that of out-of-class distribution. We further justify the use of holistic guidance by the following. One could learn a separation between classes in the feature space using a triplet or contrastive loss on an occluded dataset alone. Yet, the model may overfit on the occluded dataset due to class overlap. However, holistic datasets are much larger than occluded ones. Therefore, they can provide a good generalization of non-corrupted DCD given that class overlap can be well dealt within the dissimilarity space. By matching the DCD, the student network can learn the class overlap of the teacher network (which is well separated).

Fig. 2 shows our overall architecture along with the backbone auto-encoder-based deep feature extractor. We simultaneously take two input images—one from the holistic dataset and the other from the occluded dataset. The extracted deep features are simultaneously optimized for identity loss by learning a set of two fully connected layers for classification. We use two separate classifiers, one for the teacher to learn holistic data identities and the other for the student to learn occluded data identities.

Attention Embedding is particular to student network alone capable of producing attention for the partial features such that the attended partial feature will have good separation between within and between-class distances similar to that of the teacher. Let the attention embedding (which is a set of two layers  $1 \times 1$  convolutional filters with ReLU between them followed by batch normalization layer with sigmoid activation for final attention output) be represented by  $AE$ . The attention produced by the attention embedding is given by  $A^i = AE(f_O^i)$ . The attention maps in  $A^i$  of size  $B \times C$  are obtained for each partial feature, with  $i = 1, \dots, p$  and  $p$  being the number of parts. The attention obtained is multiplied with each partial feature to obtain attended partial features,  $f_a^i = f_O^i \otimes A^i$ . The layers for occluded

image classification,  $FC_c^i$ , are applied on each of the attended partial feature  $f_a^i$ . While training on an artificially occluded dataset alone, we use a binary classifier to learn occluded or non-occluded images on artificially occluded samples similar to [64].

In order to learn the attention, the student network relies on occluded input images and distance distribution matching. Then, DCD of the occluded and holistic features are compared. Given a mini-batch of image input with occluded and holistic images  $I_O$  and  $I_N$ , partial features  $f_N^i$  and  $f_a^i$  (partial feature with attention) are extracted. We denote the class identity for the features by  $u$  and  $v$ . Therefore, for each mini-batch, we extract pairs of image features with different combinations within a batch according to:

$$d_i^{\text{wc}}(\mathbf{I}_c^u, \mathbf{I}_c^v) = \left\| P_N^{i,u} - P_N^{i,v} \right\|_2, u = v, \quad \text{and,} \quad d_i^{\text{bc}}(\mathbf{I}_c^u, \mathbf{I}_c^v) = \left\| P_N^{i,u} - P_N^{i,v} \right\|_2, u \neq v. \quad (4)$$

Eqn. 4 transforms the features to *dissimilarity space*.  $P^i$  denotes the part features i.e.,  $f_N^i$  for holistic data and  $f_a^i$  for occluded data. The distance distributions are extracted from  $d_i^{\text{wc}}$  and  $d_i^{\text{bc}}$  for both holistic and occluded data. We implicitly learn to produce a good attention embedding  $AE$  by minimising the discrepancy between DCD of holistic and occluded data using MMD [43]. Let  $\mathbf{D}_c^{\text{wc}}$  and  $\mathbf{D}_c^{\text{bc}}$  be the distributions from  $d_i^{\text{wc}}$  and  $d_i^{\text{bc}}$ . The losses measuring the discrepancy between class distributions of holistic and occluded data are,

$$\mathcal{L}_D^{\text{wc}} = \text{MMD}(\mathbf{D}_N^{\text{wc}}, \mathbf{D}_O^{\text{wc}}), \quad \mathcal{L}_D^{\text{bc}} = \text{MMD}(\mathbf{D}_N^{\text{bc}}, \mathbf{D}_O^{\text{bc}}) \quad \text{and,} \quad \mathcal{L}_{\text{global}} = \text{MMD}(f_N, f_a) \quad (5)$$

It is important to note that losses  $\mathcal{L}_D^{\text{wc}}$  and  $\mathcal{L}_D^{\text{bc}}$  are optimized by fixing  $\mathbf{D}_N^{\text{bc}}$  and  $\mathbf{D}_N^{\text{wc}}$ , the DCD of teacher network. This allows the student network distance distribution to match that of the teacher network. Hence, Eqn. 5 calculates the discrepancies between the within-class and between-class DCDs of holistic datasets and occluded datasets. This loss is minimized during learning to obtain a good attention map from the embedding to focus on non-occluded regions of occluded images.  $\mathcal{L}_{\text{global}}$  calculates the MMD distance between teacher features and student features to encourage the model to perform well on both occluded and holistic data. This  $\mathcal{L}_{\text{global}}$  is particularly used when the holistic and occluded datasets are from different domains. Parameters  $\lambda_1$ ,  $\lambda_2$ , and  $\lambda_3$  balance these losses, and are determined empirically.

$$\mathcal{L}_D = \lambda_1 \mathcal{L}_D^{\text{wc}} + \lambda_2 \mathcal{L}_D^{\text{bc}} + \lambda_3 \mathcal{L}_{\text{global}}. \quad (6)$$

**End-to-End Learning and Testing:** The full system is optimized for reconstruction losses and identity losses with cross-entropy on both occluded and holistic data, and finally, the class distance distribution loss. The overall loss function  $\mathcal{L}_{\text{Total}}$  is given by,

$$\mathcal{L}_{\text{Total}} = \alpha \mathcal{L}_{\text{joint},c} + (1 - \alpha) \mathcal{L}_D, \quad (7)$$

where  $\alpha$  balances these losses, and is determined empirically (see supp. material). During testing, only the student is used along with the components denoted within Fig. 2 to extract  $\psi$ , which is a concatenation of global and local features,  $F_O$  and  $f_a^i$ , for both gallery and query images. Extracted features are matched with Euclidean or Cosine distance to retrieve the identity of a query image from the gallery.

## 4 Experimental Results and Discussion

**Datasets:** Our approach is validated on three challenging groups of datasets – Holistic ReID, Occluded-ReID, and Partial ReID datasets. Our main objective is to address performance on

Occluded-ReID problems (Occluded-DukeMTMC [26] and Occluded-ReID [50] datasets) and Partial ReID problems (Partial-ILIDS [26, 46] and Partial-ReID [47] datasets), but we also evaluated on Holistic problems – Market1501 [42] and Duke-MTMC [29] datasets – to further assess the effectiveness our approach on regular ReID problems. The Occluded-DukeMTMC dataset [26] contains a total of 15,618 training and 17,661 gallery with 2,210 occluded query images. This is a subset of the Duke-MTMC dataset. To test on Occluded-Duke, we train the student model with Occluded-Duke training set similar to [26, 56, 50]. The Occluded-ReID dataset [50] mimic real-world application scenarios by collecting datasets using mobile camera equipment on campus. It has a total of 2,000 annotated images with 200 identities. Each identity consists of 5 full-body images and 5 partial images. The Holistic ReID and Partial ReID datasets are described in the supplementary material.

**Implementation Details:** For validation, the ResNet50 [42] was implemented as our backbone Encoder. Transposed convolution layers were used along with Interpolation for the Decoder (see details in the supplementary material). To evaluate Occluded-Duke-MTMC, we train the student using the train data of Occluded-Duke-MTMC, and the teacher with Market1501 [42], as in the SOA. Since Occluded-ReID and Partial ReID do not have a prescribed set of training images, the whole dataset was used for testing (as in [42]). To have a common setting with the SOA [26, 27], we use an input image size of 384 X 128. We train our backbone with Partial Features with  $p = 6$  parts, and set the co-efficient  $\lambda = 0.01$ . The teacher network is pre-trained for 15 epochs, and the student-teacher is trained together for 120 epochs. The Adam optimization is used with an initial learning rate of 0.0003. We report the Cumulative Matching Characteristics (CMC) and mean average precision (mAP) [42].

**Results with Occluded and Partial ReID Problems.** Tabs. 1 and 2 show the result of our method on the Occluded-Duke and Occluded-ReID dataset compared with State-Of-The art (SOA) methods. [26, 56, 50] are Occluded-ReID methods. We show the results for our method with both ResNet50 and ResNet-IBN [27] backbones. In the Table "PM"-pose maps, "KP"-key -point detection. Our method outperforms all the other Occluded Person ReID methods mentioned in the table. Our proposed HG method performs competitively in Rank-1 without any external input, such as pose or segmentation masks, on the Occluded-Duke dataset and 2.5% on the Occluded-ReID dataset. Since the Occluded-ReID dataset does not contain training images, we show two sets of results, one with student trained on artificially occluded Market1501 dataset -HG(Unsup) and the other the student model trained on occluded samples from Occluded-Duke dataset -HG(Sup). Our results show that student trained on augmented Market1501 already outperforms many SOTA. Additionally, when the student is trained on a totally different occluded dataset (Occluded-Duke) from that of the

Method	Backbone	Supervision	Accuracy			
			Rank-1	Rank-5	Rank-10	mAP
LOMO+XQDA [45], CVPR 2015	-	None	8.1	17.0	22.0	5.0
Part Aligned [46], ICCV 2017	GoogLeNet	None	28.8	44.6	51.0	20.2
Random Erasing [47], AAAI 2020	ResNet50	None	40.5	59.6	66.8	30.0
HACNN [48], CVPR 2018	Custom	None	34.4	51.9	59.4	26.0
Adver Occluded [49], CVPR 2018	ResNet50	None	44.5	-	-	32.3
PCB [50], ECCV 2018	ResNet50	None	42.6	57.1	62.9	33.7
Part Bilinear [51], ECCV 2018	GoogLeNet	Occluded-Duke	36.9	-	-	-
PGFA [52], ICCV, 2019	ResNet50	Occluded-Duke + PM	51.4	68.6	74.9	37.3
Depth Occln [53], PRL 2020	ResNet50	Occluded-Duke + PM	53.0	67.0	72.9	38.1
HOREID [54], CVPR 2020	ResNet50	Occluded-Duke + KP	55.1	-	-	43.8
PAT [55], CVPR 2021	ResNet50	Occluded-Duke	64.5	-	-	53.6
MOS [56], AAAI 2021	ResNet50	Occluded-Duke	61.0	-	-	49.2
MOS [56], AAAI 2021	ResNet50-IBN	Occluded-Duke	<b>66.6</b>	-	-	<b>55.1</b>
HG(ours)	ResNet50	Occluded-Duke	61.4	77.0	79.8	50.5
HG(ours)	ResNet50-IBN	Occluded-Duke	65.1	<b>79.1</b>	<b>81.4</b>	54.7

Table 1: Accuracy of HG and state-of-the-art methods on the Occluded-Duke dataset.



Method	Backbone	Supervision	Accuracy			
			Rank-1	Rank-5	Rank-10	mAP
IDE [1], ICCV 2015	-	None	52.6	68.7	76.6	46.4
MLFN [8], CVPR 2018	Custom	None	42.3	60.6	68.5	36.0
HACNN [2], CVPR 2018	Custom	None	29.1	44.7	54.7	26.1
PCB [3], ECCV 2018	ResNet50	None	59.3	75.2	83.2	53.2
Part Bilinear [4], ECCV 2018	GoogLeNet	None	54.9	70.8	77.7	50.3
teacher-S [5], ArXiv 2019	ResNet	Split Test Set	55.0	64.5	77.3	59.8
PGFA [6], ICCV 2019	ResNet50	PM	57.1	77.9	84.0	56.2
PVPM [7], CVPR 2020	ResNet50	PM	70.4	84.1	89.8	61.2
HOReID [9], CVPR 2020	ResNet50	KP	80.3	-	-	70.2
PAT [10]	ResNet50	-	81.6	-	-	72.1
HG (ours Unsup)	ResNet50	None	79.4	88.5	93.7	71.1
HG (ours Sup)	ResNet50	Occluded-Duke	82.3	89.7	94.1	71.7
HG (ours Sup)	ResNet50-IBN	Occluded-Duke	<b>82.8</b>	<b>90.1</b>	<b>94.6</b>	72.0

Table 2: Accuracy of HG and state-of-the-art methods on the Occluded-ReID dataset.

test dataset (Occluded-ReID) still outperforms all SOTA. We also evaluate our method on partial ReID datasets [11] and the result table is presented in the supplementary material. We can observe from the results that our method has improved over the SOTA by 2%.

**Ablation Study.** An extensive study was conducted on Occluded-ReID to analyze the impact on training data and architecture performance. Results are shown in Tab. 3.

**(a) Impact of Data.** This study was performed on the full student-teacher model with all the components. In Tab. 3, "Augm. on Holistic": refers to student model trained with artificially occluded Market1501 dataset by Random Erasing [12]. We also use additional "occluded or non-occluded" binary classification on the features and show results with and without this classification. We can see that augmentation has helped the student model to learn Occluded-ReID and perform better than many SOTA. Using the Occluded-Duke to train the student shows that the student model performs even better on the Occluded-ReID dataset. Also, Occluded-Duke and Occluded-ReID datasets have no overlap.

**(b) Impact of Architecture.** To analyze the effect of architecture on the result from above, the teacher alone (pre-trained on Market1501) is fine-tuned on Occluded-Duke. Results indicate that fine-tuning the baseline network on Occluded-Duke performs poorly on Occluded-ReID. In Tab 3, the remarks "identity+reconstruction" refer to the experiment where student-teacher model has been trained with identity loss and reconstruction loss alone. The reconstruction loss has helped the student to learn some generative properties, and hence the overall result on Occluded-ReID data is better than the baseline fine-tuned on the Occluded-Duke dataset. We further extend our experiments by using distribution loss, identity loss along with reconstruction loss, but no attention "reconstruction+MMD." The "MMD+attention" experiment does not use reconstruction loss, and hence the autoencoder is not trained. Finally, we train the student model with the full system "reconstruction+MMD+attention," which includes the attention mechanism learned by matching distributions.

Model	Experiment	Training Data (Student)	Rank-1
<b>Impact of Training Architecture (Test Set: Occluded-ReID)</b>			
HG (student+teacher)	w/o occlusion classifier	Augment Holistic	75.6
	w/ occlusion classifier	Augment Holistic	79.4
	w/o occlusion classifier	Occluded-Duke	<b>82.3</b>
<b>Impact of Losses (Test Set: Occluded-ReID)</b>			
teacher only	identity only	-	59.5
teacher+student	identity only	Occluded-Duke	63.8
teacher+student	identity+reconstruction	Occluded-Duke	67.7
teacher+student	identity+reconstruction+MMD	Occluded-Duke	76.1
teacher+student	MMD + attention (no autoencoder)	Occluded-Duke	80.1
teacher+student	reconstruction+MMD+attention	Occluded-Duke	<b>82.3</b>
<b>Impact of backbone (Only ResNet50 Backbone, w/o Part Pooling)</b>			
teacher only	identity	-	51.3
teacher+student	reconstruction+MMD+attention	Occluded-Duke	68.5

Table 3: Impact on HG accuracy of training data and architecture on Occluded-ReID dataset.

Model	Backbone	Market1501		Duke-MTMC	
		Rank-1	mAP	Rank-1	mAP
PCB [24], ECCV 2018	ResNet50	92.3	77.4	81.8	66.1
MaskReID [25], CVPR 2018	MSCAN	90.0	75.3	-	-
FPR [26], ICCV 2019	ResNet50	95.4	<b>86.6</b>	88.6	<b>78.4</b>
PVPM [27], CVPR 2020	ResNet50	93.0	80.8	83.6	72.6
PGFA [28], ICCV 2019	ResNet50	91.2	76.8	82.6	65.5
HOReID [29], CVPR 2020	ResNet50	94.2	84.9	86.9	75.6
MOS [30], CVPR 2021	ResNet50-IBN	95.4	89.0	<b>90.6</b>	<b>80.2</b>
PAT [31], CVPR 2021	ResNet50	95.4	88.0	88.8	78.2
HG (Ours)	ResNet50	<b>95.6</b>	86.1	87.1	77.5

Table 4: Accuracy of our HG method on Market1501 and Duke-MTMC datasets.

**(c) Impact of CNN Backbone.** In order to assess the impact of part pooling, we perform another the "Impact of backbone" experiment, where only a ResNet-50 backbone is used with no part pooling. Although the overall results are this case is lower than when using part pooled features, it can be seen that proposed loss has still improves Occluded-ReID performance. From the results, we can conclude that our system learns occluded-ReID by using either artificially-occluded examples or real-occluded examples with holistic data as a reference. We show additional ablation studies in the supplementary materials.

**(d) Holistic ReID Datasets.** Tab. 4 shows results with the HG model on holistic Market1501 [44] and Duke-MTMC-[49] datasets. Results show that HG is competitive with other SOTA models designed to work with the Occluded-ReID problem.




Figure 3: Activation maps generated for four occluded images of the Partial ReID dataset by the PGFA method [28] versus our HG method. PGFA uses pose estimation additionally.

**Qualitative Results.** Fig. 3 compares activation maps from the student model tested on examples with Partial-ReID dataset with activation maps of [28] (uses pose maps for attention). In the figure, both head and legs are occluded. But, from the activation maps of our proposed HG method, it can be seen that our method is good at localising non-occluded regions alone.

## 5 Conclusion

In this paper, a novel HG student-teacher model is proposed for occluded person ReID that only requires image identity labels but no costly process to focus on visible parts of occluded regions. The proposed HG teacher considers the DCD of among samples in a holistic dataset to train a student model to generate attention maps, thereby alleviating the occlusion problem. Unlike most methods in the literature that use external supervision (such as pose) to generate visibility cues, we only rely upon the distribution of holistic data during training, using it as a soft label. Hence during test time, our model requires no external cues such as pose, and the overall parameters include just the backbone Encoder and a small embedding to generate attention maps to be used during feature extraction. Joint learning of a denoising autoencoder was used to improve the ability to self-recover from occlusion. Results on several challenging datasets show that our HG method can outperform state-of-the-art models for Occluded-ReID, as well as Holistic ReID tasks.

**Acknowledgments:** This research was supported by the Compute Canada and MITACS.

## References

- [1] E. Ahmed, M. Jones, and T.K. Marks. An improved deep learning architecture for person re-identification. In CVPR, 2015.
- [2] H. Cai, Z. Wang, and J. Cheng. Multi-scale body-part mask guided attention for person re-identification. In CVPR Workshops, 2019.
- [3] X. Chang, T. M Hospedales, and T. Xiang. Multi-level factorisation net for person re-identification. In CVPR, 2018.
- [4] W. Chen, X. Chen, J. Zhang, and K. Huang. Beyond triplet loss: a deep quadruplet network for person re-identification. In CVPR, 2017.
- [5] D. Cheng, Y. Gong, S. Zhou, J. Wang, and N. Zheng. Person re-identification by multi-channel parts-based cnn with improved triplet loss function. In CVPR, 2016.
- [6] L. Cheng, J. Wang, Y. Gong, and Q. Hou. Robust deep auto-encoder for occluded face recognition. In ACMM, 2015.
- [7] Y.MG Costa, D. Bertolini, A. S Britto, G. DC Cavalcanti, and L. ES Oliveira. The dissimilarity approach: a review. Artificial Intelligence Review, 53(4):2783–2808, 2020.
- [8] R. PW Duin, M. Bicego, M. Orozco-Alzate, S.-W. Kim, and M. Loog. Metric learning in dissimilarity space for improved nearest neighbor performance. In IAPR Workshop, 2014.
- [9] G. Ekladios, H. Lemoine, E. Granger, K. Kamali, and S. Moudache. Dual-triplet metric learning for unsupervised domain adaptation in video face recognition. In IJCNN, 2020.
- [10] E. Elhamifar, G. Sapiro, and R. Vidal. Finding exemplars from pairwise dissimilarities via simultaneous sparse recovery. NeurIPS, 2012.
- [11] M. Farenzena, L. Bazzani, A. Perina, V. Murino, and M. Cristani. Person re-identification by symmetry-driven accumulation of local features. In CVPR, 2010.
- [12] S. Gao, J. Wang, H. Lu, and Z. Liu. Pose-guided visible part matching for occluded person reid. In CVPR, 2020.
- [13] A. Gretton, K.M. Borgwardt, M.J. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample test. JMLR, 13(1):723–773, 2012.
- [14] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In CVPR, 2016.
- [15] L. He, J. Liang, H. Li, and Z. Sun. Deep spatial feature reconstruction for partial person re-identification: Alignment-free approach. In CVPR, 2018.
- [16] L. He, Y. Wang, W. Liu, H. Zhao, Z. Sun, and J. Feng. Foreground-aware pyramid reconstruction for alignment-free occluded person re-identification. In ICCV, 2019.
- [17] H. Huang, D. Li, Z. Zhang, X. Chen, and K. Huang. Adversarially occluded samples for person re-identification. In CVPR, 2018.

- [18] D. W Jacobs, D. Weinshall, and Y. Gdalyahu. Classification with non-metric distances: Image retrieval and class representation. IEEE Transactions on Pattern Analysis and Machine Intelligence, 22(6):583–600, 2000.
- [19] M. Jia, X. Cheng, Y. Zhai, S. Lu, S. Ma, Y. Tian, and J. Zhang. Matching on sets: Conquer occluded person re-identification without alignment. In AAAI, 2021.
- [20] M.M. Kalayeh, E. Basaran, M. Gökmen, M. E Kamasak, and M. Shah. Human semantic parsing for person re-identification. In CVPR, 2018.
- [21] D. Li, X. Chen, Z. Zhang, and K. Huang. Learning deep context-aware features over body and latent parts for person re-identification. In CVPR, 2017.
- [22] W. Li, X. Zhu, and S. Gong. Harmonious attention network for person re-identification. In CVPR, 2018.
- [23] Yulin Li, Jianfeng He, Tianzhu Zhang, Xiang Liu, Yongdong Zhang, and Feng Wu. Diverse part discovery: Occluded person re-identification with part-aware transformer. In CVPR, 2021.
- [24] S. Liao, Y. Hu, X. Zhu, and S.Z. Li. Person re-identification by local maximal occurrence representation and metric learning. In CVPR, 2015.
- [25] H. Liu, J. Feng, M. Qi, J. Jiang, and S. Yan. End-to-end comparative attention networks for person re-identification. IEEE Transactions on Image Processing, 26(7):3492–3506, 2017.
- [26] J. Miao, Y. Wu, P. Liu, Y. Ding, and Y. Yang. Pose-guided feature alignment for occluded person re-identification. In ICCV, 2019.
- [27] Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at once: Enhancing learning and generalization capacities via ibn-net. In ECCV, 2018.
- [28] R. Panda, A. Bhuiyan, V. Murino, and A. K. Roy-Chowdhury. Unsupervised adaptive re-identification in open world dynamic camera networks. In CVPR, 2017.
- [29] E. Ristani, F. Solera, R. Zou, R. Cucchiara, and C. Tomasi. Performance measures and a data set for multi-target, multi-camera tracking. In ECCV workshop, 2016.
- [30] C. Song, Y. Huang, W. Ouyang, and L. Wang. Mask-guided contrastive attention model for person re-identification. In CVPR, 2018.
- [31] C. Su, J. Li, S. Zhang, J. Xing, W. Gao, and Q. Tian. Pose-driven deep convolutional model for person re-identification. In ICCV, 2017.
- [32] Y. Suh, J. Wang, S. Tang, T. Mei, and K.M. Lee. Part-aligned bilinear representations for person re-identification. ECCV, 2018.
- [33] Y. Sun, L. Zheng, Y. Yang, Q. Tian, and S. Wang. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In ECCV, 2018.
- [34] H. Tan, X. Liu, S. Tian, B. Yin, and X. Li. Mhsa-net: Multi-head self-attention network for occluded person re-identification. CoRR, abs/2008.04015, 2020.

- [35] R.R. Varior, M. Haloi, and G. Wang. Gated siamese convolutional neural network architecture for human re-identification. In ECCV, 2016.
- [36] G. Wang, S. Yang, H. Liu, Z. Wang, Y. Yang, S. Wang, G. Yu, E. Zhou, and J. Sun. High-order information matters: Learning relation and topology for occluded person re-identification. In CVPR, 2020.
- [37] L. Wei, S. Zhang, H. Yao, W. Gao, and Q. Tian. Glad: Global-local-alignment descriptor for pedestrian retrieval. In ACM, 2017.
- [38] D. Wu, K. Zhang, S.-J. Zheng, Y.-T. Hao, F.-Q. Liu, X. Qin, F. Cheng, Y. Zhao, Q. Liu, C.-A. Yuan, et al. Random occlusion recovery for person re-identification. Journal of Imaging Science and Technology, 63(3):30405–1, 2019.
- [39] D. Yi, Z. Lei, S. Liao, and S.Z. Li. Deep metric learning for person re-identification. In ICPR, 2014.
- [40] L. Zhang, T. Xiang, and S. Gong. Learning a discriminative null space for person re-identification. In CVPR, 2016.
- [41] S. Zhang, L. Wen, X. Bian, Z. Lei, and S.Z. Li. Occlusion-aware r-cnn: detecting pedestrians in a crowd. In ECCV, 2018.
- [42] X. Zhang, H. Luo, X. Fan, W. Xiang, Y. Sun, Q. Xiao, W. Jiang, C. Zhang, and J. Sun. Alignedreid: Surpassing human-level performance in person re-identification. CoRR, abs/1711.08184, 2017.
- [43] L. Zhao, X. Li, Y. Zhuang, and J. Wang. Deeply-learned part-aligned representations for person re-identification. In ICCV, 2017.
- [44] L. Zheng, L. Shen, L. Tian, S. Wang, J. Wang, and Q. Tian. Scalable person re-identification: A benchmark. In ICCV, 2015.
- [45] L. Zheng, Y. Huang, H. Lu, and Y. Yang. Pose-invariant embedding for deep person re-identification. IEEE Transactions on Image Processing, 28(9), 2019.
- [46] W.-S. Zheng, S. Gong, and T. Xiang. Person re-identification by probabilistic relative distance comparison. In CVPR, 2011.
- [47] W.-S. Zheng, X. Li, T. Xiang, S. Liao, J. Lai, and S. Gong. Partial person re-identification. In ICCV, 2015.
- [48] Z. Zheng, X. Yang, Z. Yu, L. Zheng, Y. Yang, and J. Kautz. Joint discriminative and generative learning for person re-identification. In CVPR, 2019.
- [49] Z. Zhong, L. Zheng, G. Kang, S. Li, and Y. Yang. Random erasing data augmentation. In AAAI, 2020.
- [50] S. Zhou, J. Wu, F. Zhang, and P. Sehdev. Depth occlusion perception feature analysis for person re-identification. Pattern Recognition Letters, 138:617–623, 2020.
- [51] J. Zhuo, Z. Chen, J. Lai, and G. Wang. Occluded person re-identification. In ICME, 2018.
- [52] J. Zhuo, J. Lai, and P. Chen. A novel teacher-student learning framework for occluded person re-identification. CoRR, abs/1907.03253, 2019.