



# Occlusion-Guided Feature Purification Learning via Reinforced Knowledge Distillation for Occluded Person Re-Identification

Yufei ZHENG<sup>1</sup>, Yong WU<sup>2</sup>, Jiawei LIU<sup>1,3,†</sup>, Wenjun WANG<sup>1</sup>, Wenjun GAN<sup>1</sup>

<sup>1</sup> School of Information Science and Technology, University of Science and Technology of China, Hefei 230027, China

<sup>2</sup> AI Lab, China Merchants Bank, Chengdu 610041, China

<sup>3</sup> School of Computing, National University of Singapore, 119077, Singapore

<sup>†</sup>E-mail: jwliu6@ustc.edu.cn

Received: September 11, 2025 / Revised: November 2, 2025 / Accepted: November 11, 2025 / Published online: November 14, 2025

**Abstract:** Occluded person re-identification aims to retrieve holistic images of a given identity based on occluded person images. Most existing approaches primarily focus on aligning visible body parts using prior information, applying occlusion augmentation to predefined regions, or complementing the missing semantics of occluded body parts with the assistance of holistic images. Nevertheless, they struggle to generalize across diverse occlusion scenarios that are absent from the training data and often overlook the pervasive issue of feature contamination caused by holistic images. In this work, we propose a novel Occlusion-Guided Feature Purification Learning via Reinforced Knowledge Distillation (OGFR) to address these two issues simultaneously. OGFR adopts a teacher-student distillation architecture that effectively incorporates diverse occlusion patterns into feature representation while transferring the purified discriminative holistic knowledge from the holistic to the occluded branch through reinforced knowledge distillation. Specifically, an Occlusion-Aware Vision Transformer is designed to leverage learnable occlusion pattern embeddings to explicitly model such diverse occlusion types, thereby guiding occlusion-aware robust feature representation. Moreover, we devise a Feature Erasing and Purification Module within the holistic branch, in which an agent is employed to identify low-quality patch tokens of holistic images that contain noisy negative information via deep reinforcement learning, and substitute these patch tokens with learnable embedding tokens to avoid feature contamination and further excavate identity-related discriminative clues. Afterward, with the assistance of knowledge distillation, the student branch effectively absorbs the purified holistic knowledge to precisely learn robust representation regardless of the interference of occlusions. Extensive experiments validate OGFR, on Occluded-Duke it achieves 76.6% Rank-1 and 64.7% mAP, outperforming the closest Transformer-based method by +3.3% Rank-1 and +2.4% mAP, with consistent gains on other benchmarks.

**Keywords:** Occluded person re-identification; holistic person re-identification; feature purification; reinforcement learning; knowledge distillation

<https://doi.org/10.64509/jicn.12.31>

## 1 Introduction

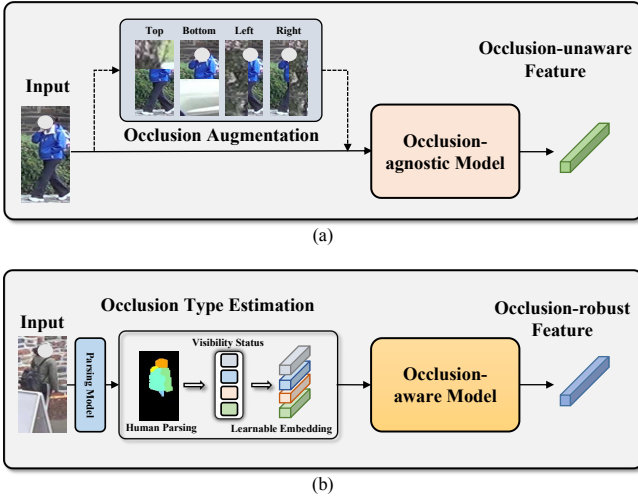
Person Re-Identification (Re-ID) is the task of identifying the same pedestrian across disjoint camera views, which has extensive applications in video surveillance, smart retail, and automated tracking *etc.* In recent years, significant advancements in person Re-ID [1–5] have led to impressive performance improvements. However, most existing methods focus on holistic visible images, while overlooking diverse

occlusions encountered in real-world scenarios, such as those caused by vehicles, trees, and crowds. Therefore, it is essential to explore effective approaches for occluded person Re-ID. Compared with holistic person Re-ID, occluded person Re-ID is more challenging due to body regions covered by obstacles and other pedestrians, which drastically reduce the availability of discriminative visual cues in feature representation.

<sup>†</sup> Corresponding author: Jiawei Liu

\* Academic Editor: Xueyang Fu

© 2025 The authors. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).



**Figure 1:** Comparison of OGFR with existing occluded person ReID methods. (a) Existing occlusion-agnostic models either completely neglect occlusion types or rely on occlusion augmentation strategy to passively generate additional occluded images, which are inherently constrained to addressing a limited range of simple occlusions. (b) OGFR leverages an additional parsing model to obtain parsing information for Occlusion Type Estimation and proactively incorporates the resulting occlusion type cues into learnable occlusion pattern embeddings, enabling effective handling of complex occlusion scenarios.

To handle above-mentioned challenges, some part-to-part matching methods employ pre-defined rules, *e.g.*, stripe partition strategy [6, 7] or introduce external models, such as human parsing [8] and pose estimation [9–11] to learn visible part representations from non-occluded regions. These methods attempt to mitigate occlusion effects by focusing exclusively on visible body parts. However, restricting attention solely to non-occluded regions overlooks the impact of different occlusion types on the overall image representation. As a result, the learned representations become overly dependent on the specific visible regions, leading to degraded performance when faced with diverse occlusions. An alternative line of research [12, 13] explores occlusion augmentation techniques to simulate occlusion conditions by artificially generating occluded images. However, these methods typically apply occlusion augmentation to fixed spatial regions, such as occluding the top, bottom, left, or right parts of an image. This results in only four basic occlusion types, which fail to capture the complexity and variability of real-world occlusions. Consequently, these methods exhibit limited generalization capabilities and struggle under diverse and unpredictable occlusion scenarios. Moreover, these methods provide only the visual content of the image as input, without explicitly incorporating occlusion-specific information. This omission forces models to infer occlusion patterns purely from visual features, without access to additional contextual cues that could facilitate a more precise understanding. As illustrated in Figure 1(a), models that either disregard occlusion types entirely or apply fixed-type occlusion augmentations, both of which lack to incorporate occlusion-related cues in the input, are collectively referred to as occlusion-agnostic models. Such models lack the ability to

distinguish between different occlusion types explicitly, making them unstable in feature representation learning and prone to poor generalization across diverse occlusion scenarios. Alongside the aforementioned approaches, feature recovery methods [14, 15] which leverage holistic image and visible near-neighbor to complement occluded body parts, have garnered increasing attention in recent years. While these methods aim to restore missing information and improve feature integrity, they still face several limitations. First, the extracted holistic and visible near-neighbor features inherently contain cluttered background and irrelevant semantics, resulting in prominent feature contamination issue. Second, the recovered representations may remain incomplete and semantically ambiguous due to severely occluded scenarios.

Motivated by the above observations, we propose a novel Occlusion-Guided Feature Purification Learning via Reinforced Knowledge Distillation (OGFR), a teacher-student distillation framework that copes with diverse occlusion scenarios by learning occlusion-aware representation, and effectively excavates and transfers purified discriminative holistic knowledge from the holistic branch to the occluded branch through reinforced knowledge distillation for mitigating feature contamination. Specifically, as shown in Figure 2, OGFR consists of a teacher branch for holistic person images and a student branch for occluded ones, equipped with Occlusion-Aware Vision Transformer (OA-ViT) and Feature Erasing and Purification (FEP) module. As illustrated in Figure 1(b), OA-ViT facilitates precise and diverse occlusion type estimation under arbitrary occlusions, eliminating the constraints of predefined occlusion regions. This allows for more flexible occlusion augmentations, significantly improving the model’s generalization ability. Additionally, OA-ViT encodes different occlusion types into learnable embeddings, incorporating explicit occlusion-related knowledge into the model. By enriching the input with both the original image and its corresponding occlusion conditions, OA-ViT transforms the model into an “occlusion-aware” framework, allowing it to explicitly aware of the occlusion status for developing a deeper understanding of occlusion patterns and learning robust features even in complex occlusion scenarios. Given that the features are incomplete when occlusion occurs, relying solely on occluded images is insufficient for comprehensive feature recovery. To address this, we introduce knowledge distillation, where occlusion-augmented data is input into the student branch to model occluded scenarios, while unaltered data is fed into the teacher branch to capture the holistic information. Through distillation, the student branch is guided to learn complementary features from the teacher branch, enhancing its ability to reconstruct identity-related representations. Considering the inherently existing cluttered background and useless semantics in holistic images, we do not directly transfer the complementary knowledge from the holistic branch to the occluded branch, instead designing additional FEP module within the holistic branch. Unlike indiscriminately removing all tokens belonging to the background, FEP selectively eliminates noise and irrelevant information while retaining potentially valuable information. This is achieved through a deep reinforcement learning-based erasing mask generation agent, which dynamically identifies and filters out

low-quality patch tokens in holistic images, using the performance improvement over the case without the binary mask as the reward signal. By eliminating these low-quality tokens that negatively impact the model's performance, the agent improves feature purity and overall effectiveness. To further refine representation learning, FEP replaces erased patch features with learnable embedding tokens, leveraging a customized two-layer Transformer decoder to extract potential information and purify discriminative cues. After that, with the aid of knowledge distillation, the occluded branch assimilates such purified complementary knowledge to learn identity-related robust representation. Extensive experiments on the popular occluded and holistic ReID datasets demonstrate the effectiveness and superiority of OGFR.

The main contributions can be summarized as follows:

- We propose a novel Occlusion-Guided Feature Purification Learning via Reinforced Knowledge Distillation (OGFR) approach for occluded person Re-ID. This teacher-student distillation framework effectively integrates diverse occlusion types into feature representation and facilitates the transfer of purified holistic knowledge from the holistic branch to the occluded branch.
- We introduce an Occlusion-Aware Vision Transformer that encodes visibility status of different body parts into learnable occlusion embeddings to guide occlusion-aware robust feature representation, ensuring adaptability to various occlusion scenarios.
- We design a Feature Erasing and Purification module within the holistic branch to identify low-quality patches through deep reinforcement learning and replace them with learnable embedding tokens for excavating refined identity-related information from holistic person images.
- Our OGFR achieves state-of-the-art performance on several occluded and holistic ReID benchmarks, including Occluded-Duke [16], Occluded-reID [17], P-DukeMTMC-reID [17], Market-1501 [18], DukeMTMC-reID [19] and MSMT17 [20].

## 2 Related Work

### 2.1 Holistic Person Re-Identification

Over the past decade, a variety of holistic person Re-ID methods have sprung up, aiming at extracting discriminative and robust representations from full-body pedestrian images. Early studies [21–24] primarily focused on designing hand-crafted features with body structures or robust distance metric learning. For example, Hirzer *et al.* [23] employed metric learning to identify a suitable space for matching samples from different cameras. Instead of utilizing traditional Euclidean distance metrics, they adopted discriminative Mahalanobis distance measures, leading to a significant enhancement in performance. Recently, deep learning technique has been widely adopted for learning discriminative deep representations of holistic pedestrians. Sun *et al.* [1] proposed a part-based convolutional baseline for learning local features, which leverages a simple stripe partitioning strategy and refined part pooling to enhance intra-part consistency. Wang *et al.* [25] proposed the Multiple Granularity Network (MGN),

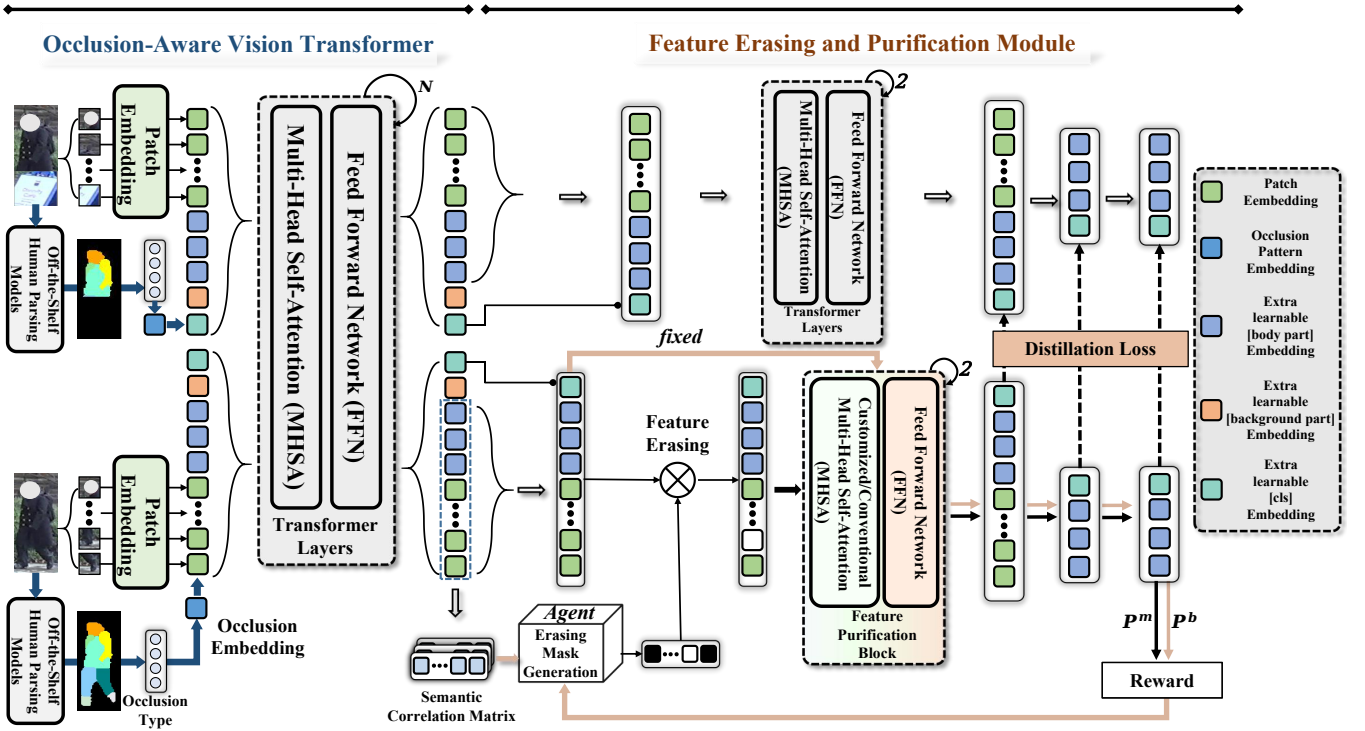
which consists of one branch for global feature representations and two branches for local feature representations, aiming at integrating discriminative information with various granularities. Zheng *et al.* [26] developed a coarse-to-fine pyramid model via a multi-loss dynamic training scheme to incorporate local and global discriminative information. More recently, He *et al.* [2] pioneered the use of pure transformer framework for object Re-ID, and proposed two novel modules to enhance the capacity of feature learning.

### 2.2 Occluded Person Re-Identification

Compared to holistic person Re-ID, occluded person Re-ID presents greater challenges. Existing approaches can be roughly divided into three categories: part-to-part matching methods, occlusion augmentation methods, and feature recovery methods. Part-to-part matching methods exploit predefined rules or external models to extract and align visible part representations from non-occluded regions to mitigate the occlusion effect. For instance, Zhu *et al.* [27] clustered feature map pixels as pseudo-labels for body parts, enabling visible part-based representation learning. However, these methods focus only on visible regions, completely ignoring the impact of different occlusion types. To enhance robustness against various occlusion patterns, some studies [13] introduced occlusion augmentation, artificially generating occluded images to simulate real-world scenarios. Chen *et al.* [12] apply occlusions at four fixed positions and enforce occlusion-type score prediction to handle occlusions. However, these techniques typically apply occlusions to fixed spatial regions and rely solely on raw visual content, forcing the model to infer occlusions implicitly, which restricts its adaptability to different complex occlusion scenarios. Beyond occlusion augmentation, feature recovery methods seek to reconstruct missing information by leveraging visible near-neighbor features and holistic image representations. Hou *et al.* [28] proposed a Region Feature Completion block that captures long-range spatial and long-term temporal contexts for feature recovery. Xu *et al.* [15] designed a feature recovery transformer that leverages k-nearest neighbor features to restore occluded representations. However, these approaches risk feature contamination, as holistic and near-neighbor features often include background noise and irrelevant semantics. Different from these existing methods, our proposed OGFR not only actively and adaptively handles various occlusion types but also simultaneously avoids feature contamination in holistic images.

### 2.3 Knowledge Distillation

Knowledge Distillation is a model compression technique that transfers knowledge from a high-capacity teacher model to a lightweight student model, enabling parameter reduction while maintaining performance. This technique is initially proposed by Hinton *et al.* [29], the fundamental idea is to guide the training of the student model using soft labels (probability distributions) generated by the teacher model, rather than relying solely on hard labels (single class labels). This enables the student model to acquire deep-level knowledge and enhance generalization capabilities from the teacher model, leading to improved performance. Several studies



**Figure 2:** The overview of Occlusion-Guided Feature Purification Learning via Reinforced Knowledge Distillation approach. It contains a teacher branch for holistic person images and a student branch for occluded ones with OA-ViT and FEP modules.

[30–34] have applied knowledge distillation to person Re-ID. For example, Zheng *et al.* [30] utilized knowledge distillation to transfer knowledge from the pose-guided branch to the main branch during the training, which enables the discarding of pose-based branches while maintaining satisfied performance. In our approach, knowledge distillation is employed to address the challenges of occluded person Re-ID. We devised a teacher-student distillation framework where the teacher model operates on holistic images, while the student model focuses on occlusion scenarios. Through reinforced knowledge distillation, we effectively transfer purified, discriminative knowledge from the holistic branch to the occluded branch, enabling the student model to better handle occlusion scenarios.

### 3 Proposed Method

In this section, we provide a detailed introduction to our proposed method, Occlusion-Guided Feature Purification Learning via Reinforced Knowledge Distillation (OGFR). As depicted in Figure 2, OGFR comprises a dual-branch architecture, integrating a teacher branch and a student branch. The teacher branch integrates a weight-shared Occlusion-Aware Vision Transformer (OA-ViT) and a customized Feature Erasing and Purification (FEP) module, while the student branch consists of the same weight-shared Occlusion-Aware Vision Transformer along with two conventional transformer layers.

#### 3.1 Occlusion-Aware Vision Transformer

In order to adaptively handle various occlusion scenarios, OA-ViT reformulates different occlusion types as the set of visibility status of body parts through an occlusion label estimator, and incorporates occlusion type clues into learnable

occlusion embeddings to facilitate occlusion-aware robust feature learning via occlusion-aware vision encoder.

##### 3.1.1 Occlusion Label Estimator

Occlusion occurs when certain body parts of a pedestrian are obstructed by external objects, leading to partial loss of body integrity. Based on this, different occlusion types in pedestrian images can be uniformly represented as a set of visibility statuses of various body parts. Given a person image, we first generate parsing labels  $M \in \mathbb{R}^{H \times W \times (K+1)}$  with the pre-trained MASK-RCNN [35] and PifPaf [36] models, where  $H, W$  denote the height and width of the image, and  $K+1$  is the defined  $K$  semantic body parts and background (For instance,  $K=8$  defines the following semantic body parts: *head*, *left/right arm*, *torso*, *left/right leg* and *left/right feet*). The element  $M[h, w, k] \in \{0, 1\}$  is the parsing result at spatial position  $[h, w]$ , indicating whether belonging to the  $k$ -th body part or to the background. Once we obtain the parsing labels  $M$  from the pre-trained model, we assign occlusion types based on the number of pixels within each body part as indicated by these labels.

However, due to domain discrepancies, the segmentation results from the parsing model are not always precise. To mitigate this issue and obtain a more stable occlusion type representation, we further merge these  $K$  fine-grained semantic body parts into  $C$  coarse-grained body parts ( $K > C$ ). The corresponding coarse semantic parsing labels  $\tilde{M} \in \mathbb{R}^{H \times W \times C}$  (excluding the background) are generated by the merging operation. It can be observed that if a body part is occluded by an obstacle, the integrity of that part is compromised and fewer pixels in the image belong to that body part. Considering that, we calculate the total number of pixels corresponding to each body part to determine its visibility status, *i.e.*, *Occluded* or



*Non-Occluded*. It is formulated as follows:

$$\mathbf{y}(c) = \sum_{h=0}^{H-1} \sum_{w=0}^{W-1} \widetilde{\mathbf{M}}[h, w, c] \quad (1)$$

$$l_c = \begin{cases} 1 & \text{if } \mathbf{y}(c) < \lambda, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

where  $\mathbf{y}(c)$  denotes the total number of pixels belong to  $c$ -th coarse-grained body part,  $\lambda$  is a pre-defined threshold, and  $l_c$  denotes the visibility status of  $c$ -th coarse-grained body part (1 for *Occluded* and 0 for *Non-Occluded*). We finally acquire the occlusion type representation for arbitrary images,  $\mathbf{z} = [l_0, l_1, \dots, l_{C-1}] \in \mathbb{R}^C$ .

### 3.1.2 Occlusion-Aware Vision Encoder

We incorporate the obtained occlusion type information of the image into occlusion pattern embeddings, and insert them into Vision Transformer (ViT) [37] together with different types of embeddings, endowing the model with the capacity to perceive and handle diverse occlusion scenarios. Specifically, we initialize the learnable occlusion pattern embeddings as  $\mathbf{O}_e \in \mathbb{R}^{N_0 \times \dots \times N_{C-1} \times D}$ , where  $\{N_0, \dots, N_{C-1}\}$  denotes the 2-dimension visibility status for each body part. Therefore, the occlusion pattern embeddings  $\mathbf{E}_o \in \mathbb{R}^D$  for occlusion type  $\mathbf{z}$  can be formulated as below:

$$\mathbf{E}_o = \mathbf{O}_e[\mathbf{z}[0], \dots, \mathbf{z}[C-1]] \quad (3)$$

In addition to the occlusion pattern embeddings, other essential types of embeddings, such as position embeddings and camera embeddings, are also appended to the input sequence. These provide crucial location and camera-specific information for the patch tokens in the sequence. As for patch tokens, we first split the image with a resolution of  $H \times W$  into  $N$  patches based on the following formulation:

$$N = \left\lfloor \frac{H+S-P}{S} \right\rfloor \times \left\lfloor \frac{W+S-P}{S} \right\rfloor \quad (4)$$

where  $S$  and  $P$  denote stride and patch size,  $\lfloor \cdot \rfloor$  represents the floor function. Upon obtaining patch tokens  $\mathbf{P}_s \in \mathbb{R}^{N \times D}$  through the projection operation, we introduce a learnable [cls] token  $x_{cls}$  to capture the global representation. Moreover,  $K+1$  learnable part tokens  $\mathbf{P}_l \in \mathbb{R}^{(K+1) \times D}$  are appended to the input sequence for the adaptive learning of part representations, which are semantically aligned with  $K$  fine-grained body parts and the background. Notably, we maintain  $K$  fine-grained body parts segmentation instead of  $C$  coarse-grained body parts segmentation to enhance the discriminative capabilities of these part representations. Finally, the input sequence is described as follows:

$$\mathbf{E} = [\mathbf{x}_{cls}; \mathbf{P}_l; \mathbf{P}_s] + \mathbf{P}_E + \gamma_1 \mathbf{E}_o + \gamma_2 \mathbf{C}_e \quad (5)$$


where  $\mathbf{x}_{cls}$ ,  $\mathbf{P}_l$  and  $\mathbf{P}_s$  are concatenated to form the main body of the sequence. Note that the position of the learnable tokens  $\mathbf{P}_l$  can be changed, *e.g.*, by placing them right after the patch tokens or shuffling their internal order. For simplicity,

we adopt the case where they are placed on the left. Additionally,  $\mathbf{P}_E$ ,  $\mathbf{E}_o$ , and  $\mathbf{C}_e$  are directly added to these tokens. Here,  $\mathbf{P}_E$  denotes position embeddings,  $\mathbf{C}_e$  represents camera embeddings proposed by [2].  $\gamma_1$  and  $\gamma_2$  are hyper-parameters used to balance the contributions of these embedding components. We feed the input sequence into the occlusion-aware Vision Transformer to obtain the robust representation  $\mathbf{F} \in \mathbb{R}^{(N+K+2) \times D}$ , consisting of  $\{\mathbf{f}^g, \mathbf{f}^0, \mathbf{f}^1, \dots, \mathbf{f}^K, \mathbf{f}^p\}$ , where  $\mathbf{f}^g$  denotes the global feature,  $\mathbf{f}^p \in \mathbb{R}^{N \times D}$  denotes the patch features,  $\mathbf{f}^0$  denotes the part feature corresponding to the background, and  $\{\mathbf{f}^1, \dots, \mathbf{f}^K\}$  represent the part features corresponding to different body parts.

Based on the occlusion-aware ViT, we establish a teacher-student architecture, where the branch for holistic person images serves as the teacher and the branch for occluded ones acts as the student. Within this architecture, we extract  $\mathbf{F}_h$  and  $\mathbf{F}_o$  through the weight-shared occlusion-aware ViT, where  $h, o$  denote the holistic and occluded branches, respectively. The holistic feature  $\mathbf{F}_h$  is then refined via a Feature Erasing and Purification module to produce the purified feature  $\mathbf{F}_{hp}$ . In parallel, the occluded feature  $\mathbf{F}_o$  is processed through two conventional transformer layers, resulting in the enhanced feature  $\mathbf{F}_{op}$ . These features,  $\mathbf{F}_{hp}$  and  $\mathbf{F}_{op}$ , form the foundation for the subsequent knowledge distillation process.

## 3.2 Feature Erasing and Purification Module

To explore meaningful discriminated clues while mitigating interference from irrelevant noise, we design the FEP Module within the holistic branch to excavate refined identity-related information in holistic images. As illustrated in Figure 3, FEP module trains an erasing mask generation agent that identifies low-quality patch tokens from all patch tokens and then substitutes these patch features with learnable embeddings to obtain purified discriminative representations via a customized two-layer decoder.



**Figure 3:** The detailed structure of the Feature Erasing and Purification Module.

### 3.2.1 Erasing Mask Generation

Given that the Markov Decision Process (MDP) provides a formal framework for modeling decision-making by defining key components essential for reinforcement learning, we formalize the process of erasing mask generation as a one-step MDP. Reinforcement learning is employed to enable the erasing mask generation agent to adaptively filter low-quality patch tokens based on the performance gains relative to the case without using the erasing mask. By integrating these

elements, the framework ensures an efficient and dynamic approach to improving model performance. The MDP is composed of state, agent, action, and reward, and is described as follows: *State*: The semantic correlation matrix  $\mathbf{I} \in \mathbb{R}^{N \times (K+1)}$  is defined as follows:

$$\mathbf{I} = \text{norm}(\mathbf{f}^p) \cdot \text{norm}([\mathbf{f}^0, \dots, \mathbf{f}^K])^T \quad (6)$$

which encodes the affiliation relationship between different image patches and a specific body part, serves as the state representation. Here,  $\text{norm}(\cdot)$  denotes row-wise L2 normalization. It is calculated by evaluating the similarity between patch features  $\mathbf{f}_h^p$  and part features  $\{\mathbf{f}_h^k\}_{k=0}^K$ . *Agent*: The agent generates an erasing mask using a fully connected (FC) layer followed by a sigmoid activation function, in which each element indicates the probability that a specific patch token is retained. It can be described as follows:

$$\mathbf{m} = \text{sigmoid}(\mathbf{I}\mathbf{G}^T) \quad (7)$$

where  $\mathbf{G} \in \mathbb{R}^{1 \times (K+1)}$  represents linear transformation. *Action*: The action determines whether each patch is selected, with values of 1 (to *retain*) or 0 (to *replace*). The decision follows a Bernoulli distribution, introducing stochasticity to prevent the agent from converging to suboptimal local solutions:

$$\mathbf{a} = \text{Bernoulli}(\mathbf{m}) \quad (8)$$

*Reward*: We devise a reward mechanism that provides incentives for using the erasing mask whose predicted scores for the true identity labels exceed those of the baseline. We follow the steps below to obtain baseline predictions. Without using the erasing mask, we directly feed the representation  $\{\mathbf{f}_h^g, \mathbf{f}_h^1, \dots, \mathbf{f}_h^K, \mathbf{f}_h^p\}$  into the customized two-layer decoder. Then the classifier outputs Re-ID prediction scores. We compute the average accuracy of predicted scores  $p^b$  as the baseline. During the training of the agent, we keep the parameters of the baseline fixed and update its value every 10 epochs. At each update, we refresh  $p^b$  via an online moving average as follows:

$$p^b = \frac{p^b \times n + p_{new}^b}{n + 1} \quad (9)$$

where  $n$  is the number of updates so far and  $p_{new}^b$  is the latest accuracy of predicted scores. Similarly, following the procedure above, we represent the predicted scores with using the erasing mask as  $p^m$ . The reward is defined as follows:

$$r = \frac{p^m - p^b}{1 - p^b} \quad (10)$$

At each time step  $t$  within an episode, the agent observes the state  $s_t$ , takes an action  $a_t$  to generate the erasing mask  $p^m$ , and receives an immediate reward  $r(s_t, a_t)$ . The total reward for the  $b$ -th episode is given by  $R_b = \sum_t r(s_t, a_t)$ . We employ the policy gradient method [38] to find the optimal set of parameters  $\theta^*$  representing the policy function, aiming to maximize the expected cumulative reward  $\pi_\theta$ . We define

the optimization objective as  $J(\pi_\theta)$ . To approximate the gradient, we execute the agent for  $B$  episodes and subsequently compute the average gradient as:

$$\nabla_\theta J(\theta) \approx \frac{1}{B} \sum_{b=1}^B \left( \sum_{t=1}^N \nabla_\theta \log \pi_\theta(a_{b,t} | s_{b,t}) \right) R_b \quad (11)$$

where  $a_{b,t}$  is the action taken by the agent for the erasing mask in  $b$ -th episode.

### 3.2.2 Feature Purification

We determine whether to retain or replace the patch tokens based on the action  $\mathbf{a}$ . When the action is 0, the corresponding low-quality patch tokens (patch features) are replaced with learnable embedding tokens. This replacement mechanism effectively discards redundant semantics and suppresses noise. After obtaining the substituted patch features with the remaining unchanged ones, we concatenate them with global and part features, which is denoted as  $\tilde{\mathbf{F}}_h$ . This refined representation is subsequently fed into the customized two-layer decoder, formulated as follows:

$$\begin{aligned} \mathbf{F}'_h &= \text{FFN}(\text{LN}(\text{MHSA}_{\text{fix}}(\tilde{\mathbf{F}}_h))) \\ \mathbf{F}_{hp} &= \text{FFN}(\text{LN}(\text{MHSA}(\mathbf{F}'_h))) \end{aligned} \quad (12)$$

where FFN and LN denote a feed-forward layer and a layer normalization layer, respectively.  $\text{MHSA}_{\text{fix}}$  is a specialized multi-head self-attention layer, which maintains the global and part features fixed and only allows the patch features learnable to thoroughly excavate potential identity-related clues. In  $\text{MHSA}_{\text{fix}}$ , each head independently processes the input features and derives the queries, keys, and values as follows:

$$\mathbf{Q} = W_q \tilde{\mathbf{f}}_h^p \quad \mathbf{K} = W_k \tilde{\mathbf{F}}_h \quad \mathbf{V} = W_v \tilde{\mathbf{F}}_h \quad (13)$$

where  $W_q, W_k, W_v$  are weights of linear projections,  $\tilde{\mathbf{f}}_h^p$  represents the patch features in  $\tilde{\mathbf{F}}_h$ , which are used to compute the queries  $\mathbf{Q}$ . Through the attention mechanism, attention weights for  $\mathbf{Q}$  relative to the keys  $\mathbf{K}$  are computed, and these weights are utilized to aggregate  $\mathbf{V}$ , enabling the update of patch features:

$$\mathbf{F}'_{ha} = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{k}}\right)\mathbf{V} \quad (14)$$

The fixed global and part features can be obtained by directly taking their corresponding values from  $\mathbf{V}$ , which can be represented by the following formula:

$$\mathbf{F}'_{ha}{}^{\text{fixed}} = \{\mathbf{v}^g, \mathbf{v}^1, \dots, \mathbf{v}^K\} \quad (15)$$

Subsequently, the updated patch features are concatenated with the fixed global and part features to construct the complete representation.

$$\mathbf{F}'_{ha} = [\mathbf{F}'_{ha}{}^{\text{fixed}}, \mathbf{F}'_{ha}{}^p] \quad (16)$$

Following the above steps, each attention head produces an output feature  $F'_{ha}$ . The outputs from all heads are then concatenated and passed through a linear transformation to obtain the final representation of MHSA<sub>fix</sub>. This is followed by LN and FFN, ultimately yielding the final refined representation  $F'_h$ . Next, we apply MHSA, which denotes a conventional multi-head self-attention layer. This mechanism facilitates mutual learning among the global feature, the semantic part features, and the patch features, allowing them to interact and refine each other. By integrating MHSA, LN and FFN, we obtain the refined purified discriminative feature  $F_{hp}$ , which contains  $\{f_{hp}^g, f_{hp}^1, \dots, f_{hp}^K, f_{hp}^p\}$ .

### 3.3 Knowledge Distillation

To fully leverage identity-related clues from the holistic branch, we employ knowledge distillation to transfer purified identity-related knowledge to the occluded branch, enabling the occluded branch to assimilate holistic knowledge and effectively handle occlusion scenarios. First of all, by feeding the features  $\{f_o^g, f_o^1, \dots, f_o^K, f_o^p\}$  into the two transformer layers, we obtain the features  $F_{op}$  for the occlusion branch, which is composed of  $\{f_{op}^g, f_{op}^1, \dots, f_{op}^K, f_{op}^p\}$ . Then we employ the following loss functions to promote the transfer of knowledge.

(1) The mean square error loss is utilized to minimize the Euclidean distance between  $F_{hp}$  and  $F_{op}$  by:

$$\mathcal{L}_{mse} = \|F_{hp} - F_{op}\|_2^2 \quad (17)$$

(2) We employ the cosine similarity loss to minimize the discrepancies in high-level global and part representations of two branches for enforcing feature consistency, which is formulated as:

$$\mathcal{L}_{cos} = \frac{1}{K+1} \sum_i \left(1 - \frac{f_{hp}^i \cdot f_{op}^i}{\|f_{hp}^i\| \cdot \|f_{op}^i\|}\right), i \in \{g, 1, \dots, K\} \quad (18)$$

(3) We incorporate the conventional knowledge distillation loss to further strengthen the knowledge transfer process, which is defined as follows:

$$\mathcal{L}_{kd} = \frac{1}{K+1} \sum_i (\alpha \mathcal{L}_{id}(f_{op}^i) + \beta \text{KL}(f_{hp}^i \| f_{op}^i)) \quad (19)$$

where  $i \in g, 1, \dots, K$ , and the identity loss  $\mathcal{L}_{id}$  denotes cross-entropy loss that enforces alignment between the predicted logits, derived from the feature representation  $f_{op}^i$ , and ground truth identity labels [12]. KL denotes Kullback-Leibler Divergence.  $\alpha$  and  $\beta$  are hyper-parameters to balance the loss terms.

### 3.4 Model Optimization

Firstly, we employ the identity loss to optimize the global features for excavating more discriminative fine-grained clues, which can be formulated as:

$$\mathcal{L}_{en} = \mathcal{L}_{id}(f_o^g) + \mathcal{L}_{id}(f_h^g) \quad (20)$$

In addition to the aforementioned loss, we impose supervision on the semantic correlation matrix using parsing labels  $M$  in section 3.1.1 as the ground-truth. This is achieved via a cross-entropy loss, formulated as follows:

$$\mathcal{L}_{mask} = - \sum_{k=0}^K \sum_{h=0}^{H-1} \sum_{w=0}^{W-1} M_k(h, w) \log(\tilde{I}_k(h, w)) \quad (21)$$

where  $\tilde{I} \in \mathbb{R}^{H \times W \times (K+1)}$  denotes the transformed semantic correlation matrix obtained by performing rearrangement and bilinear interpolation operations on the matrix  $I$ .  $\tilde{I}_k(h, w)$  denotes the predicted probability for part  $k$  at spatial location  $(h, w)$ . Moreover, we integrate the triplet loss [39] to encourage the holistic branch to learn more discriminative purified features, as formulated below:

$$\mathcal{L}_{tr} = \mathcal{L}_{tri}(f_{hp}^g) + \frac{1}{K} \sum_{i=1}^K \mathcal{L}_{tri}(f_{hp}^i) \quad (22)$$

The overall objective function for optimization during the training stage is defined as follows:

$$\mathcal{L} = (\mathcal{L}_{mse} + \mu_1 \mathcal{L}_{cos} + \mathcal{L}_{kd}) + (\mathcal{L}_{en} + \mu_2 \mathcal{L}_{tr}) + \mathcal{L}_{mask} \quad (23)$$

where  $\mu_1$  and  $\mu_2$  are hyper-parameters to adjust the corresponding loss terms.

During the inference stage, we employ the occluded branch to extract global and part features for calculating the similarity scores. Specifically, we only consider the similarity of visible features. The visibility label for the global feature is set to 1. For part features, visibility is assessed by examining pixel values in the transformed semantic correlation matrix  $\tilde{I}$ . A part is considered visible if the maximum pixel value within its corresponding region exceeds 0, formulated as follows:

$$v_i = \begin{cases} 1 & \text{if } \max_{h,w} \tilde{I}_i(h, w) > 0 \\ 0 & \text{otherwise.} \end{cases} \quad (24)$$

where  $i \in \{1, 2, \dots, K\}$ . Inspired by [16], the distance between visibility-based query images and gallery images is computed as follows:

$$dist = \frac{\sum_{i=1}^K (v_i^q \cdot v_i^g) \cdot D(f_{op}^{i,q}, f_{op}^{i,g}) + D(f_{op}^{g,q}, f_{op}^{g,g})}{\sum_{i=1}^K (v_i^q \cdot v_i^g) + 1} \quad (25)$$

where  $D(\cdot, \cdot)$  refers to the Euclidean distance.

## 4 Experiments

In this section, we conduct extensive experiments on occluded and holistic datasets to verify the superiority of our proposed method. Then, we report ablation studies to analyze the effectiveness of each key component. Finally, we further provide the visualization results to intuitively show the effectiveness of our proposed method.

## 4.1 Datasets and Evaluation Metrics

For experimental evaluation, we conduct experiments on six public person ReID benchmarks, including three occluded datasets (Occluded-Duke[16], Occluded-reID [17], P-DukeMTMC-reID [17]) and three holistic datasets (Market-1501 [18], DukeMTMC-reID [19] and MSMT17 [20]). **Note:** The DukeMTMC-reID datasets have been abandoned by their proposers due to privacy concerns. In compliance with ethical and legal requirements, we report experimental results on related datasets for reference only, as these datasets are widely used for this task, including DukeMTMC-reID, Occluded-Duke, and P-DukeMTMC-reID, and we do not provide any processing code or trained model parameters derived from them.

**Occluded-Duke** [16] is an occluded Re-ID dataset containing 35,489 images from 8 cameras, with 15,618 for training and 2,210/17,661 for query/gallery. **Occluded-reID** [17] is a small-scale occluded dataset with 2,000 images of 200 identities, where each identity has five full-body and five occluded images. **P-DukeMTMC-reID** [17] is a occluded Re-ID dataset modified from DukeMTMC-reID [19], comprising 24,143 images of 1,299 identities. The training set includes 12,927 images, while the testing set consists of 2,163 query images and 9,053 gallery images. **Market-1501** [18] is a holistic Re-ID dataset with 12,936 training images and 3,368/15,913 for query/gallery, captured by 6 cameras. **DukeMTMC-reID** [19] is a challenging holistic dataset containing 16,522 training images of 702 identities, with 2,228/17,661 for query/gallery. **MSMT17** [20] is a large-scale holistic dataset with 126,441 images of 4,101 identities from 15 cameras, including 32,621 for training and 11,659/82,161 for query/gallery.

**Evaluation Metrics.** We adopt widely used Cumulative Matching Characteristic (CMC) at Rank- $k$  and mean Average Precision (mAP) as evaluation metrics.

## 4.2 Implementation Details

We adopt ViT [37] model pre-trained on ImageNet [40] as our backbone. Concretely, all input images are resized to  $256 \times 128$ . During training, we apply data augmentation techniques, including random horizontal flipping, padding pixels, random cropping and occlusion augmentation. For occlusion augmentation, we adopt the techniques, including random horizontal flipping, padding pixels, random cropping and occlusion augmentation. For occlusion augmentation, we adopt the technique of selectively superimposing an obstacle patch onto a randomly chosen region within a pedestrian image, following previous methods [41, 42]. We utilize SGD as the optimizer with the weight decay of  $1e-4$ . The learning rate is initiated as 0.008 with cosine learning rate decay, which is decayed at 40 and 70 epochs. The batch size is set to 64 with 4 images per identity. Following TransReID [2], we adopt a smaller stride and set  $\gamma_1, \gamma_2$  to 3.0. The dimension of the features  $D$  is set to 768. The number of  $K$  is set to 8. The number of  $C$  is 4 containing the coarse-grained semantic regions {head, arms, torso, legs}. The hyper-parameter  $\lambda$  is set to 5. The  $\alpha$  and  $\beta$  are set to 0.3 and 0.4. The  $\mu_1$  and  $\mu_2$  are set to 0.5. We implement our framework under two Nvidia Tesla v100 GPUs.

**Table 1:** Performance comparison with state-of-the-art methods on Occluded-Duke and Occluded-REID datasets. The first group: CNN-based methods; the second group: Transformer-based methods.

Method	References	Occluded-Duke		Occluded-reID	
		Rank-1	mAP	Rank-1	mAP
PVPM [10]	CVPR 20	-	-	70.4	61.2
MHSA [43]	TNNLS 22	59.7	44.8	-	-
Pirt [44]	ACM MM 21	60.0	50.9	-	-
IGOAS [45]	TIP 21	60.1	49.4	81.1	-
RTGAT [46]	TIP 23	61.0	50.1	71.8	51.0
VGTri [7]	ICCV 21	62.2	46.3	81.0	71.0
OAMN [12]	ICCV 21	62.6	46.1	-	-
ISP [27]	ECCV 20	62.8	52.3	-	-
PGFL-KD [30]	ACM MM 21	63.0	54.1	80.7	70.3
RFC [28]	TPAMI 21	63.9	54.5	-	-
QPM [47]	TMM 23	64.4	49.7	-	-
OPR-DAAO [42]	TIFS 22	66.2	55.4	84.2	75.1
MoS [48]	AAAI 21	66.6	55.1	-	-
PAT [49]	CVPR 21	64.5	53.6	81.6	72.1
DRL-Net [41]	TMM 22	65.8	53.9	-	-
TransReID [2]	ICCV 21	66.4	59.2	-	-
FED [13]	CVPR 22	68.1	56.4	86.3	79.3
AMG [50]	TMM 23	68.5	59.7	-	-
SAP [34]	AAAI 23	70.0	62.2	83.0	76.8
FRT [15]	TIP 22	70.7	61.3	80.4	71.0
DPM [51]	ACM MM 22	71.4	61.8	85.5	79.7
OAT [52]	TIP 24	71.8	62.2	82.6	78.2
MAHATMA [53]	TCSVT 25	73.3	62.3	85.8	79.5
<b>OGFR</b>		<b>76.6</b>	<b>64.7</b>	<b>86.9</b>	<b>80.9</b>

**Table 2:** Performance comparison with state-of-the-art methods on P-DukeMTMC-reID dataset. The first group: CNN-based methods; the second group: Transformer-based methods.

Method	References	P-DukeMTMC-reID			
		Rank-1	Rank-5	Rank-10	mAP
MHSA [43]	TNNLS 22	70.7	81.0	84.6	41.1
PCB [1]	ECCV 18	79.4	87.1	90.0	63.9
PVPM [10]	CVPR 20	85.1	91.3	93.3	69.9
RTGAT [46]	TIP 23	85.6	91.5	93.4	74.3
PGFA [16]	ICCV 19	85.7	92.0	94.2	72.4
ISP [27]	ECCV 20	89.0	94.1	95.3	74.7
QPM [47]	TMM 23	90.7	94.4	95.9	75.3
FED* [13]	CVPR 22	90.9	94.4	95.4	79.4
DPM* [51]	ACM MM 22	91.5	95.1	96.0	82.5
TransReID* [2]	ICCV 21	92.4	95.8	96.9	82.6
<b>OGFR</b>		<b>93.2</b>	<b>96.2</b>	<b>97.1</b>	<b>83.4</b>

\* indicates that the results are reproduced from the official code.

## 4.3 Comparison with State-of-the-art Methods

### 4.3.1 Results on Occluded Datasets

To validate the effectiveness of our proposed OGFR, we conduct a comprehensive comparison with state-of-the-art methods across three occluded Re-ID datasets: Occluded-Duke, Occluded-REID, and P-DukeMTMC-reID. As shown in Table 1 and Table 2, the compared methods are categorized based on their architectural design into CNN-based and



**Table 3:** Performance comparison with state-of-the-art methods on three holistic datasets including Market-1501, DukeMTMC-reID and MSMT17. The first group: CNN-based methods; the second group: Transformer-based methods. Hybrid denotes ResNet50 + transformer.

Method	References	Backbone	Market-1501		DukeMTMC-reID		MSMT17	
			Rank-1	mAP	Rank-1	mAP	Rank-1	mAP
OAMN [12]	ICCV 21	ResNet50	93.2	79.8	86.3	72.6	-	-
IGOAS [45]	TIP 21	ResNet50	93.4	84.1	86.9	75.1	-	-
Pirt [44]	ACM MM 21	ResNet50-ibn	94.1	86.3	88.9	77.6	-	-
MHSA [43]	TNNLS 22	ResNet50	94.6	84.0	87.3	73.1	-	-
OPR-DAAO [42]	TIFS 22	ResNet50	95.1	86.2	88.5	76.5	-	-
RFC [28]	TPAMI 21	ResNet50	95.2	89.2	90.7	80.7	82.0	60.2
PGFL-KD [30]	ACM MM 21	ResNet50	95.3	87.2	89.6	79.5	-	-
RTGAT [46]	TIP 23	ResNet50	95.3	88.2	89.1	80.2	-	-
ISP [27]	ECCV 20	HRNet-W32	95.3	88.6	89.6	80.0	-	-
MoS [48]	AAAI 21	ResNet50-ibn	95.4	89.0	90.6	80.2	-	-
DRL-Net [41]	TMM 22	Hybrid	94.7	86.9	88.1	76.6	78.4	55.3
FED [13]	CVPR 22	ViT-B	95.0	86.3	89.4	78.0	-	-
AMG [50]	TMM 23	ViT-B	95.0	88.5	-	-	-	-
MAHATMA [53]	TCSVT 25	ViT-B	95.2	88.2	91.2	81.2	85.6	<b>68.1</b>
TransReID [2]	ICCV 21	ViT-B	95.2	88.9	90.7	82.0	83.5	64.4
PAT [49]	CVPR 21	Hybrid	95.4	88.0	88.8	78.2	-	-
FRT [15]	TIP 22	Hybrid	95.5	88.1	90.5	81.7	-	-
DPM [51]	ACM MM 22	ViT-B	95.5	89.7	91.0	82.6	-	-
OAT [52]	TIP 24	ViT-B	95.7	89.9	91.2	82.3	-	-
SAP [34]	AAAI 23	ViT-B	96.0	<b>90.5</b>	-	-	85.7	67.8
<b>OGFR</b>		ViT-B	<b>96.4</b>	90.2	<b>92.3</b>	<b>82.7</b>	<b>86.2</b>	66.9

Transformer-based approaches. Since there are no publicly available results for Transformer-based methods on the P-DukeMTMC-reID dataset, we reproduce several open-source methods to ensure a fair comparison. The CNN-based methods include PVPM [10], MHSA [43], Pirt [44], IGOAS [45], RTGAT [46], VGTri [7], OAMN [12], ISP [27], PGFL-KD [30], RFC [28], QPM [47], OPR-DAAO [42], MoS [48], PCB [1], and PGFA [16]. Meanwhile, the Transformer-based methods encompass PAT [49], DRL-Net [41], TransReID [2], FED [13], AMG [50], SAP [34], FRT [15], DPM [51], OAT [52] and MAHATMA [53]. It is evident that Transformer-based approaches consistently outperform CNN-based methods, benefiting from their superior ability to model long-range dependencies and capture complex patterns. Our OGFR follows this trend, achieving state-of-the-art performance across all datasets. Compared to the limited occlusion data-augmented method FED, OGFR achieves +8.5% Rank-1 accuracy and +8.3% mAP on Occluded-Duke, and +0.6% Rank-1 accuracy and +1.6% mAP on Occluded-REID. It also obtains an improvement of +2.3% Rank-1 accuracy and +4.0% mAP on the P-DukeMTMC-reID dataset, demonstrating its superior performance. The substantial gain is attributed to OA-ViT, which incorporates the visibility status of body parts into learnable occlusion embeddings, effectively guiding occlusion-aware feature learning for robust identity representation across varying occlusion scenarios. Additionally, compared to SAP, which also adopts a teacher-student architecture with ViT as the backbone, OGFR achieves +6.6% Rank-1 accuracy and +2.5% mAP on Occluded-Duke, and +3.9% Rank-1 accuracy and +4.1% mAP on Occluded-REID, further validating its superiority. This demonstrates the effectiveness of refined

holistic knowledge transferred from the teacher branch. Simultaneously, the student branch efficiently absorbs purified knowledge to learn identity-related discriminative features, thereby achieving outstanding performance.

#### 4.3.2 Results on Holistic Datasets

To further verify the superiority of OGFR, we conduct extensive experiments on three holistic datasets: Market-1501, DukeMTMC-reID, and MSMT17. We compare OGFR against state-of-the-art methods, which are categorized into CNN-based and Transformer-based methods, including OAMN [12], IGOAS [45], Pirt [44], MHSA [43], RFC [28], PGFL-KD [30], RTGAT [46], ISP [27], MoS [48], DRL-Net [41], FED [13], AMG [50], MAHATMA [53], OPR-DAAO [42], TransReID [2], PAT [49], FRT [15], DPM [51], OAT [52], SAP [34]. As depicted in Table 3, our proposed OGFR achieves the best Rank-1 accuracy of 96.4% and 86.2% on the Market-1501 and MSMT17 datasets, respectively. It also achieves near-optimal mAP on these datasets. Compared to the latest MAHATMA method, OGFR surpasses it by +2.0% mAP on Market-1501, demonstrating its strong capability in holistic person retrieval. Although it falls slightly behind by 1.2% mAP on the MSMT17 dataset, OGFR still maintains highly competitive performance, highlighting its robustness across diverse datasets. Furthermore, OGFR attains the highest Rank-1 accuracy of 92.3% and mAP of 82.7% on the DukeMTMC dataset. The success on these datasets demonstrates that our method not only handles occlusion effectively but also excels at capturing discriminative information in non-occluded scenarios.

## 4.4 Ablation Studies

### 4.4.1 Effectiveness of Proposed Modules

We conducted ablation studies on Occluded-Duke dataset to validate the effectiveness of each component. As shown in Table 4, The baseline is a single-branch model with occlusion augmentation but without OA-ViT, FEP, or KD, using the ViT [37] framework with identity and triplet losses. Incorporating OA-ViT improves the baseline by +2.7% in Rank-1 accuracy and +1.9% in mAP due to its ability to encode diverse occlusion patterns into occlusion pattern embeddings for robust feature learning. Adding the FEP module to the single-branch model further boosts performance by +4.0% in Rank-1 accuracy and +0.8% in mAP, as FEP adaptively filters low-quality patch tokens and replaces them with learnable embedding tokens to extract purified discriminative features. Along with the introduction of knowledge distillation (KD), the model transitions into a teacher-student structure, where the complete branch serves as the teacher and the occlusion branch as the student. Retaining FEP in the student branch improves Rank-1 accuracy by +0.6% and mAP by +1.2%, showing that the teacher provides valuable supplementary information. However, placing FEP in the teacher branch achieves better results. Further analysis shows that placing FEP exclusively in the teacher branch is optimal. The teacher branch refines holistic knowledge to obtain purer knowledge, while the student branch, without FEP, retains complete occluded image information, enabling it to leverage the rich information to learn and absorb more complete and purer knowledge from the teacher.

### 4.4.2 Analysis of the Occlusion Type in OA-ViT

We analyze the key parameter  $C$  of the occlusion type in OA-ViT. As shown in Figure 4, when the parameter  $C$  increases from 1 to 4, the model demonstrates an improvement in performance from 75.0% to 76.6% in Rank-1 accuracy. This indicates that using visibility status of finer granularity body parts allows for accurately characterize occlusion patterns, leading to superior retrieval performance. However, when increasing the parameter  $C$  from 4 to 8, OGFR achieves slight performance degradation, signifying that excessive occlusion types introduce redundant information. Thus, we set the parameter  $C$  to 4 in our experiments, which achieves optimal performance.

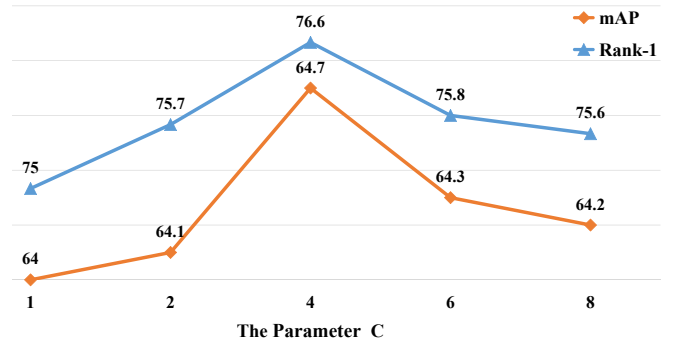
### 4.4.3 Analysis of Erasing Mask Generation in FEP

In Table 5, we compare performance under different mask generation methods in the FEP module. It can be observed that the performance of random mask generation is reduced by +0.6% mAP compared to without the mask. This suggests that random mask generation could potentially obscure crucial information. On the other hand, masking all background and occlusion tokens leads to a +0.9% improvement in Rank-1 accuracy compared to without the mask, indicating that it helps reduce some noise interference. Our erasing mask generation improves Rank-1 accuracy by +2.2% over random masking generation and +1.2% over background and occlusion masking generation. By adaptively avoiding the interference of cluttered backgrounds and irrelevant semantics

while retaining useful information, the erasing mask generation method proves to be more effective than random masking and background/occlusion masking, leading to superior performance.

**Table 4:** Ablation studies of each component in OGFR on Occluded-Duke dataset. The OA-ViT refers to the Occlusion-Aware Vision Transformer, the FEP refers to the Feature Erasing and Purification Module, -S indicates the Student branch, -T the Teacher branch, the KD refers to the Knowledge Distillation.

Method				Occluded-Duke			
OA-ViT	FEP-S	FEP-T	KD	Rank-1	Rank-5	Rank-10	mAP
-	-	-	-	65.8	81.6	86.5	58.4
✓	-	-	-	68.5	82.0	87.0	60.3
✓	✓	-	-	72.5	83.3	86.8	61.1
✓	✓	-	✓	73.1	83.7	86.8	62.3
✓	✓	✓	✓	74.3	84.8	87.3	63.6
✓	-	✓	✓	<b>76.6</b>	<b>86.4</b>	<b>89.8</b>	<b>64.7</b>




**Figure 4:** Impact of the parameter  $C$  for the occlusion type in terms of Rank-1 (%) and mAP (%) on Occluded-Duke dataset.

**Table 5:** Performance comparison with different types of masks on Occluded-Duke dataset. Here, w/o denotes without Erasing Mask Generation, the  $\mathcal{R}$  denotes Random Mask Generation, the  $\mathcal{A}$  denotes Background and Occlusion Mask Generation, the  $\mathcal{E}$  denotes our Erasing Mask Generation.

Mask Method	Occluded-Duke			
	Rank-1	Rank-5	Rank-10	mAP
w/o	74.5	86.3	89.5	64.3
$\mathcal{R}$	74.4	85.8	88.9	63.7
$\mathcal{A}$	75.4	86.2	89.6	64.3
$\mathcal{E}$	<b>76.6</b>	<b>86.4</b>	<b>89.8</b>	<b>64.7</b>

### 4.4.4 Analysis of Customized Decoder in FEP

As illustrated in Table 6, we conduct a comparative performance analysis across different types of decoders within the FEP module. It can be observed that our customized decoder exhibits a noteworthy improvement, achieving +1.3% increase in Rank-1 accuracy and +0.6% in mAP compared to the conventional decoder. The improvement can be attributed to the incorporation of a specialized multi-head self-attention layer. This layer maintains the global and part features as fixed while allowing the patch features to be learnable. Such



**Figure 5:** Performance curve of four hyper-parameters within the loss functions on Occluded-Duke dataset. (a) Performance curve of the hyperparameter  $\alpha$ , (b) Performance curve of hyperparameter  $\beta$ , (c) Performance curve of hyperparameter  $\mu_1$ , (d) Performance curve of hyperparameter  $\mu_2$ .

a design facilitates patch features in thoroughly excavating potential identity-related clues from global and semantic features. More importantly, it prevents patch features from introducing noise or contaminating the global and semantic context.

#### 4.4.5 Analysis of Hyper-parameter

To validate the robustness of hyper-parameters, we examine the impact of varying hyperparameter values on the performance of OGFR, as illustrated in Figure 5 and Table 7.

**Influence of Hyperparameter  $\alpha$ .** We evaluate the effect of hyperparameter  $\alpha$  on performance, as it plays a crucial role in balancing the knowledge distillation process. The  $\alpha$  is the coefficient of the identity loss ( $\mathcal{L}_{id}$ ) within the knowledge distillation loss ( $\mathcal{L}_{kd}$ ). As depicted in Figure 5 (a), too small  $\alpha$  results in poor performance. This is primarily because a lower  $\alpha$  weakens the preservation of crucial identity-related information during knowledge distillation. Conversely, too large  $\alpha$  may disrupt the delicate balance required for effective knowledge distillation, which degrades the performance. Therefore, in our experiments,  $\alpha$  is set to 0.3.

**Influence of Hyperparameter  $\beta$ .** We investigate the influence of hyperparameter  $\beta$ , which controls the weight assigned to the Kullback-Leibler (KL) divergence term in the knowledge distillation loss. As shown in Figure 5 (b), small

$\beta$  results in insufficient consideration of the knowledge distillation term, thereby impairing performance. While larger  $\beta$  could potentially overshadow the impact of other loss components, causing instability and degrading the performance. Thus, in our experimental setup,  $\beta$  is set to 0.4.

**Influence of Hyperparameter  $\mu_1$ .** We analyze the impact of hyperparameter  $\mu_1$  within the overall loss function, which controls the contribution of the cosine similarity loss ( $\mathcal{L}_{cos}$ ). As illustrated in Figure 5 (c), small  $\mu_1$  diminishes the significance of the cosine similarity loss, while a larger value can dominate the overall loss function, adversely affecting discriminative feature learning. Thus, we set  $\mu_1$  to 0.5 in our experimental setup.

**Influence of Hyperparameter  $\mu_2$ .** We analyze the influence of hyperparameter  $\mu_2$ , which regulates the contribution of the triplet loss ( $\mathcal{L}_{tr}$ ) within the overall loss function. As shown in Figure 5 (d),  $\mu_2 = 0.5$  achieves the best performance. This is because small  $\mu_2$  diminishes the influence of the triplet loss, potentially impairing the ability to learn discriminative features, while larger  $\mu_2$  could overly emphasize the triplet loss, potentially causing an imbalance in the learning process.

**Influence of Hyperparameter  $\lambda$ .** To assess the sensitivity of the occlusion-threshold  $\lambda$ , we conduct a controlled ablation on Occluded-Duke by varying  $\lambda \in \{0, 5, 10, 15\}$ , where a body part is deemed occluded if its human-parsing foreground pixel count is below  $\lambda$ . As shown in Table 7, performance is stable across a reasonable range, with  $\lambda = 5$  yielding the best overall results. Setting  $\lambda$  too low ( $\lambda = 0$ ) slightly increases false-visible assignments due to tiny noisy fragments, while larger thresholds ( $\lambda \geq 10$ ) introduce more false-occluded cases, mildly degrading both Rank-1 and mAP. These results validate that our method is robust to  $\lambda$  within a broad interval and that  $\lambda = 5$  strikes a favorable balance between suppressing parsing noise and preserving genuinely visible small parts.

Finally, it is worth noting that these five hyper-parameters exhibit satisfactory robustness. Across variations in their values, model performance consistently remains within a range of no more than 3%. This observation indicates that, regardless of parameter adjustments, the model maintains a relatively consistent level of robustness in terms of performance stability.

#### 4.4.6 Analysis of Knowledge Distillation

As shown in Table 8, we conduct an ablation analysis on the three losses in knowledge distillation. The performance is improved by +2.4% Rank-1 accuracy and +1.4% mAP with

**Table 6:** Performance comparison between different types of decoders on Occluded-Duke dataset. Here, the  $\mathcal{M}$  denotes a conventional decoder with two multi-head self-attention layers, while  $\mathcal{S}$  denotes our customized decoder.

Method	Occluded-Duke			
	Rank-1	Rank-5	Rank-10	mAP
$\mathcal{M}$	75.3	86.0	89.0	64.1
$\mathcal{S}$	<b>76.6</b>	<b>86.4</b>	<b>89.8</b>	<b>64.7</b>

**Table 7:** Performance comparison between different values of  $\lambda$  on Occluded-Duke dataset.

$\lambda$	Occluded-Duke			
	Rank-1	Rank-5	Rank-10	mAP
0	76.3	86.0	89.1	64.5
5	<b>76.6</b>	<b>86.4</b>	<b>89.8</b>	<b>64.7</b>
10	76.2	85.9	89.0	64.2
15	75.8	85.8	88.9	64.0

the mean squared error (MSE) loss, as the MSE loss minimizes the feature discrepancy between the holistic branch and the occluded branch. This ensures that the features of both branches are as similar as possible, thereby aligning the feature spaces of the two branches. Furthermore, compared with the combination of the MSE loss and knowledge distillation (KD) loss, the performance is improved by +1.5% Rank-1 accuracy and +1.1% mAP using the cosine similarity (COS) loss, which reduces divergence between global and semantic part representations within the two branches, promoting orientation coherence in high-level features. Finally, with the assistance of the KD loss, the performance is improved by +2.1% Rank-1 accuracy and +2.0% mAP. This is because the KD loss minimizes the difference between the probability distributions of the holistic branch and the occluded branch, encouraging the occluded branch to adhere to the distribution of the holistic branch.

**Table 8:** Ablation studies of each loss function in Knowledge Distillation on Occluded-Duke dataset.

Knowledge Distillation			Occluded-Duke			
MSE	COS	KD	Rank-1	Rank-5	Rank-10	mAP
-	✓	✓	74.2	84.9	88.5	63.3
✓	-	✓	75.1	86.0	89.0	63.6
✓	✓	-	74.5	85.9	89.4	62.7
✓	✓	✓	<b>76.6</b>	<b>86.4</b>	<b>89.8</b>	<b>64.7</b>

## 4.5 Visualization Results

As shown in Figure 6 (a), each pair of images represents the original image and the visualization of features from the occluded branch, respectively. We can observe OGFR is not affected by occlusions, allowing it to focus on more discriminative semantic clues. Moreover, we visualize the visibility status of four body parts under different occlusions in Figure 6 (b). By integrating the visibility information of body parts into the learnable occlusion pattern embeddings, OGFR effectively handles various occlusion scenarios. Furthermore, Figure 7 presents the visualization results of the erasing mask, which effectively identifies obstacles, cluttered backgrounds, and other non-target pedestrians. This is because we employ an erasing mask generation agent to dynamically identify low-quality patch tokens from all patch tokens in a deep reinforcement learning manner based on the performance gains relative to the case without employing the erasing mask.


## 4.6 Qualitative Retrieval Results

To demonstrate the effectiveness of our OGFR in handling occluded scenes, we present the comparison of Rank-5 retrieval results between the baseline and our proposed method in Figure 8, where green and red rectangles represent correct and incorrect retrieval results, respectively. For each occluded query person image on the left, where the occlusion scenarios have not appeared in the training data, our OGFR demonstrates a strong capability to accurately retrieve the target pedestrians despite the interference of various occlusion conditions, compared to the baseline. Furthermore, as observed from the first two retrieval examples in the top row,

OGFR exhibits enhanced perceptual ability in identifying target pedestrians in cases of occlusion by other individuals, resulting in superior performance. This highlights the robustness and effectiveness of OGFR in addressing real-world occlusion scenarios within pedestrian retrieval tasks.

## 4.7 Computational Complexity

As shown in Table 9, we conduct extensive experiments to demonstrate that our OGFR not only exhibits superior performance but also advantages in model parameters, floating-point operations (FLOPs), training time ( $T_T$ ), and inference time ( $T_I$ ). We compare OGFR with recent popular occluded Re-ID methods: TransReID [2], FED [13], and DPM [51], ensuring a fair comparison by using ViT-B as the backbone. All experiments are conducted on one NVIDIA GeForce GTX 1080Ti. From Table 9, it can be observed that OGFR exhibits relatively fewer model parameters while maintaining comparable FLOPs compared to the other methods. Compared to DPM [51], which does not rely on external models, our approach showcases advantages in both model parameters and FLOPs. In terms of training times, our approach, which incorporates additional modules like reinforcement learning and knowledge distillation, has a longer training time compared to TransReID. However, it remains on par with methods such as DPM and FED, which also introduce substantial architectural modifications to ViT. Importantly, during inference, only the student branch of OGFR is utilized, leading to a significant reduction in inference time. Our method exhibits a considerably lower inference latency than DPM while achieving performance comparable to TransReID. The comparative results, evaluated across multiple metrics, demonstrate that our method achieves superior performance within an appropriate level of complexity. This highlights the effectiveness and practical applicability of OGFR as a highly competitive solution for occluded Re-ID tasks.



**Figure 6:** (a) Feature visualization results of the learned representation from the occluded branch. (b) Visualization of occlusion types for four coarse-grained body parts including {head, arms, torso, legs}.

**Table 9:** Complexity comparison with state-of-the-art methods on the occluded-duke dataset. Here, M represents Million, G represents Giga, and the units for training time ( $T_T$ ) and inference time ( $T_I$ ) are both expressed in ms/image.

Method	Rank-1	mAP	Params	FLOPs	$T_T$	$T_I$
TransReID [2]	66.4	59.2	103.3M	23.6G	4.8	5.2
FED [13]	68.1	56.4	146.1M	23.2G	17.3	3.5
DPM [51]	71.4	61.8	146.1M	32.8G	12.8	14.1
<b>OGFR</b>	<b>76.6</b>	<b>64.7</b>	121.7M	26.5G	14.1	7.0





**Figure 7:** The visualization results of erasing mask in FEP. The white regions represent the areas selected by the agent for masking.



**Figure 8:** Comparison of Rank-5 retrieval results between the baseline and our proposed OGFR on the Occluded-Duke dataset. The green and red rectangles highlight positive and negative matching, respectively.

## 5 Conclusion

We presented OGFR, a unified occlusion-guided feature purification framework via reinforced knowledge distillation for occluded person re-identification. OGFR couples an Occlusion-Aware ViT that encodes diverse occlusion types into learnable embeddings with a Feature Erasing and Purification module that leverages a reinforcement learning agent to excise low-quality patch tokens and replace them with learnable embeddings, thereby mitigating feature contamination from holistic images. Through a teacher–student distillation paradigm, the student absorbs purified holistic knowledge while explicitly modeling occlusions, yielding robust, identity-discriminative representations under complex and unseen occlusion patterns. Extensive experiments on both occluded and holistic ReID benchmarks demonstrate consistent state-of-the-art performance, including notable gains on Occluded-Duke, validating the effectiveness and generalization of our design. Beyond advancing occluded ReID, our study highlights the importance of explicitly modeling

occlusion types and purifying holistic supervision, suggesting promising directions for broader vision tasks involving structured missing information and noisy context.

## Funding

This work was supported by the National Natural Science Foundation of China (NSFC) under Grant 62476260, 62225207 and 62436008, the Fundamental Research Funds for the Central Universities under Grant WK2100000057.

## Author Contributions

Conceptualization, Yufei Zheng and Jiawei Liu; methodology, Yufei Zheng and Wenjun Gan; software, Yufei Zheng, Wenjun Wang and Wenjun Gan; validation, Yufei Zheng and Wenjun Gan; formal analysis, Yufei Zheng; investigation, Yufei Zheng; resources, Jiawei Liu and Yong Wu; data curation, Yufei Zheng; writing—original draft preparation, Yufei

Zheng; writing—review and editing, Yufei Zheng; visualization, Yufei Zheng; supervision, Jiawei Liu and Yong Wu; project administration, Jiawei Liu and Yong Wu; funding acquisition, Yong Wu. All authors have read and agreed to the published version of the manuscript.

## Conflict of Interest

All the authors declare that they have no conflict of interest.

## References

- [1] Sun, Y., Zheng, L., Yang, Y., Tian, Q., Wang, S.: Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In Proceedings of the European Conference on Computer Vision, pp. 480–496 (2018). [https://doi.org/10.1007/978-3-030-01225-0\\_30](https://doi.org/10.1007/978-3-030-01225-0_30)
- [2] He, S., Luo, H., Wang, P., Wang, F., Li, H., Jiang, W.: Transreid: Transformer-based object re-identification. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 15013–15022 (2021). <https://doi.org/10.1109/ICCV48922.2021.01474>
- [3] Jin, X., Lan, C., Zeng, W., Chen, Z., Zhang, L.: Style normalization and restitution for generalizable person re-identification. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 3140–3149 (2020). <https://doi.org/10.1109/CVPR42600.2020.00321>
- [4] Yan, C., Pang, G., Jiao, J., Bai, X., Feng, X., Shen, C.: Occluded person re-identification with single-scale global representations. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 11875–11884 (2021). <https://doi.org/10.1109/ICCV48922.2021.01166>
- [5] Ren, X., Zhang, D., Bao, X., Zhang, Y.:  $s^2$ -net: Semantic and saliency attention network for person re-identification. IEEE Transactions on Multimedia **25**, 4387–4399 (2022) <https://doi.org/10.1109/TMM.2022.3174768>
- [6] Zhang, X., Yan, Y., Xue, J.-H., Hua, Y., Wang, H.: Semantic-aware occlusion-robust network for occluded person re-identification. IEEE Transactions on Circuits and Systems for Video Technology **31**(7), 2764–2778 (2020) <https://doi.org/10.1109/TCSVT.2020.3033165>
- [7] Yang, J., Zhang, J., Yu, F., Jiang, X., Zhang, M., Sun, X., Chen, Y.-C., Zheng, W.-S.: Learning to know where to see: A visibility-aware approach for occluded person re-identification. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 11885–11894 (2021)
- [8] Huang, H., Chen, X., Huang, K.: Human parsing based alignment with multi-task learning for occluded person re-identification. In IEEE International Conference on Multimedia and Expo, pp. 1–6 (2020). <https://doi.org/10.1109/ICME46284.2020.9102789>
- [9] He, L., Liu, W.: Guided saliency feature learning for person re-identification in crowded scenes. In Proceedings of the European Conference on Computer Vision, pp. 357–373 (2020)
- [10] Gao, S., Wang, J., Lu, H., Liu, Z.: Pose-guided visible part matching for occluded person reid. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 11744–11752 (2020). <https://doi.org/10.1109/ICCV.2019.00063>
- [11] Miao, J., Wu, Y., Yang, Y.: Identifying visible parts via pose estimation for occluded person re-identification. IEEE Transactions on Neural Networks and Learning Systems **33**(9), 4624–4634 (2021) <https://doi.org/10.1109/TNNLS.2021.3059515>
- [12] Chen, P., Liu, W., Dai, P., Liu, J., Ye, Q., Xu, M., Chen, Q., Ji, R.: Occlude them all: Occlusion-aware attention network for occluded person re-id. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 11833–11842 (2021)
- [13] Wang, Z., Zhu, F., Tang, S., Zhao, R., He, L., Song, J.: Feature erasing and diffusion network for occluded person re-identification. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 4754–4763 (2022)
- [14] Zhao, Y., Zhu, S., Liang, Z.: Patch features reconstruction transformer for occluded person re-identification. In 2022 41st Chinese Control Conference, pp. 6273–6278 (2022). <https://doi.org/10.23919/CCC55666.2022.9901670>
- [15] Xu, B., He, L., Liang, J., Sun, Z.: Learning feature recovery transformer for occluded person re-identification. IEEE Transactions on Image Processing **31**, 4651–4662 (2022) <https://doi.org/10.1109/TIP.2022.3186759>
- [16] Miao, J., Wu, Y., Liu, P., Ding, Y., Yang, Y.: Pose-guided feature alignment for occluded person re-identification. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 542–551 (2019)
- [17] Zhuo, J., Chen, Z., Lai, J., Wang, G.: Occluded person re-identification. In IEEE International Conference on Multimedia and Expo, pp. 1–6 (2018)
- [18] Zheng, L., Shen, L., Tian, L., Wang, S., Wang, J., Tian, Q.: Scalable person re-identification: A benchmark. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 1116–1124 (2015)

- [19] Zheng, Z., Zheng, L., Yang, Y.: Unlabeled samples generated by gan improve the person re-identification baseline in vitro. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 3754–3762 (2017)
- [20] Wei, L., Zhang, S., Gao, W., Tian, Q.: Person transfer gan to bridge domain gap for person re-identification. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 79–88 (2018)
- [21] Zheng, W.-S., Gong, S., Xiang, T.: Person re-identification by probabilistic relative distance comparison. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 649–656 (2011). <https://doi.org/10.1109/CVPR.2011.5995598>
- [22] Avraham, T., Gurvich, I., Lindenbaum, M., Markovitch, S.: Learning implicit transfer for person re-identification. In Proceedings of the European Conference on Computer Vision, pp. 381–390 (2012). [https://doi.org/10.1007/978-3-642-33863-2\\_38](https://doi.org/10.1007/978-3-642-33863-2_38)
- [23] Hirzer, M., Roth, P.M., Köstinger, M., Bischof, H.: Relaxed pairwise learned metric for person re-identification. In Proceedings of the European Conference on Computer Vision, pp. 780–793 (2012)
- [24] Chen, W., Chen, X., Zhang, J., Huang, K.: Beyond triplet loss: a deep quadruplet network for person re-identification. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 403–412 (2017). <https://doi.org/10.1109/CVPR.2017.145>
- [25] Wang, G., Yuan, Y., Chen, X., Li, J., Zhou, X.: Learning discriminative features with multiple granularities for person re-identification. In Proceedings of the ACM International Conference on Multimedia, pp. 274–282 (2018)
- [26] Zheng, F., Deng, C., Sun, X., Jiang, X., Guo, X., Yu, Z., Huang, F., Ji, R.: Pyramidal person re-identification via multi-loss dynamic training. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 8514–8522 (2019)
- [27] Zhu, K., Guo, H., Liu, Z., Tang, M., Wang, J.: Identity-guided human semantic parsing for person re-identification. In Proceedings of the European Conference on Computer Vision, pp. 346–363 (2020)
- [28] Hou, R., Ma, B., Chang, H., Gu, X., Shan, S., Chen, X.: Feature completion for occluded person re-identification. IEEE Transactions on Pattern Analysis and Machine Intelligence **44**(9), 4894–4912 (2021) <https://doi.org/10.1109/TPAMI.2021.3079910>
- [29] Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015)
- [30] Zheng, K., Lan, C., Zeng, W., Liu, J., Zhang, Z., Zha, Z.-J.: Pose-guided feature learning with knowledge distillation for occluded person re-identification. In Proceedings of the ACM International Conference on Multimedia, pp. 4537–4545 (2021)
- [31] Hu, Z., Hou, W., Liu, X.: Deep batch active learning and knowledge distillation for person re-identification. IEEE Sensors Journal **22**(14), 14347–14355 (2022) <https://doi.org/10.1109/JSEN.2022.3181238>
- [32] Ni, H., Li, Y., Gao, L., Shen, H.T., Song, J.: Part-aware transformer for generalizable person re-identification. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 11280–11289 (2023). <https://doi.org/10.1109/ICCV51070.2023.01036>
- [33] Lan, L., Teng, X., Zhang, J., Zhang, X., Tao, D.: Learning to purification for unsupervised person re-identification. IEEE Transactions on Image Processing **32**, 3338–3353 (2023) <https://doi.org/10.1109/TIP.2023.3278860>
- [34] Jia, M., Sun, Y., Zhai, Y., Cheng, X., Yang, Y., Li, Y.: Semi-attention partition for occluded person re-identification. In Proceedings of the AAAI Conference on Artificial Intelligence, pp. 998–1006 (2023)
- [35] He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 2961–2969 (2017)
- [36] Kreiss, S., Bertoni, L., Alahi, A.: Pifpaf: Composite fields for human pose estimation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 11977–11986 (2019)
- [37] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., *et al.*: An image is worth 16x16 words: Transformers for image recognition at scale. In International Conference on Learning Representations, pp. 1–218 (2021)
- [38] Peters, J.: Policy gradient methods. Scholarpedia **5**(1), 3698 (2011) <https://doi.org/10.4249/scholarpedia.3698>
- [39] Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 815–823 (2015)
- [40] Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 248–255 (2009)
- [41] Jia, M., Cheng, X., Lu, S., Zhang, J.: Learning disentangled representation implicitly via transformer for

- occluded person re-identification. *IEEE Transactions on Multimedia* **25**, 1294–1305 (2022) <https://doi.org/10.1109/TMM.2022.3141267>
- [42] Wang, S., Liu, R., Li, H., Qi, G., Yu, Z.: Occluded person re-identification via defending against attacks from obstacles. *IEEE Transactions on Information Forensics and Security* **18**, 147–161 (2022)
- [43] Tan, H., Liu, X., Yin, B., Li, X.: Mhsa-net: Multi-head self-attention network for occluded person re-identification. *IEEE Transactions on Neural Networks and Learning Systems* **32**, 1568–1582 (2023) <https://doi.org/10.1109/TNNLS.2022.3144163>
- [44] Ma, Z., Zhao, Y., Li, J.: Pose-guided inter-and intra-part relational transformer for occluded person re-identification. In *Proceedings of the ACM International Conference on Multimedia*, pp. 1487–1496 (2021). <https://doi.org/10.1145/3474085.3475283>
- [45] Zhao, C., Lv, X., Dou, S., Zhang, S., Wu, J., Wang, L.: Incremental generative occlusion adversarial suppression network for person reid. *IEEE Transactions on Image Processing* **30**, 4212–4224 (2021) <https://doi.org/10.1109/TIP.2021.3070182>
- [46] Huang, M., Hou, C., Yang, Q., Wang, Z.: Reasoning and tuning: Graph attention network for occluded person re-identification. *IEEE Transactions on Image Processing* **32**, 1568–1582 (2023) <https://doi.org/10.1109/TIP.2023.3247159>
- [47] Wang, P., Ding, C., Shao, Z., Hong, Z., Zhang, S., Tao, D.: Quality-aware part models for occluded person re-identification. *IEEE Transactions on Multimedia* **25**, 3154–3165 (2023) <https://doi.org/10.1109/TMM.2022.3156282>
- [48] Jia, M., Cheng, X., Zhai, Y., Lu, S., Ma, S., Tian, Y., Zhang, J.: Matching on sets: Conquer occluded person re-identification without alignment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 1673–1681 (2021)
- [49] Li, Y., He, J., Zhang, T., Liu, X., Zhang, Y., Wu, F.: Diverse part discovery: Occluded person re-identification with part-aware transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2898–2907 (2021). <https://doi.org/10.1109/CVPR46437.2021.00292>
- [50] Mao, J., Yao, Y., Sun, Z., Huang, X., Shen, F., Shen, H.-T.: Attention map guided transformer pruning for occluded person re-identification on edge device. *IEEE Transactions on Multimedia* **25**, 1592–1599 (2023) <https://doi.org/10.1109/TMM.2023.3265159>
- [51] Tan, L., Dai, P., Ji, R., Wu, Y.: Dynamic prototype mask for occluded person re-identification. In *Proceedings of the ACM International Conference on Multimedia*, pp. 531–540 (2022)
- [52] Li, Y., Liu, Y., Zhang, H., Zhao, C., Wei, Z., Miao, D.: Occlusion-aware transformer with second-order attention for person re-identification. *IEEE Transactions on Image Processing* **33**, 3200–3211 (2024) <https://doi.org/10.1109/TIP.2024.3393360>
- [53] Zhang, G., Yang, Y., Zheng, Y., Martin, G., Wang, R.: Mask-aware hierarchical aggregation transformer for occluded person re-identification. *IEEE Transactions on Circuits and Systems for Video Technology* **35**(6), 5821–5832 (2025) <https://doi.org/10.1109/TCSVT.2025.3531142>