

Dynamic Feature Pruning and Consolidation for Occluded Person Re-identification

YuTeng Ye¹, Hang Zhou¹, Jiale Cai¹, Chenxing Gao¹, Youjia Zhang¹, Junle Wang², Qiang Hu³, Junqing Yu¹, Wei Yang^{1*}

¹Huazhong University of Science and Technology, Wuhan, China

²Tencent

³Shanghai Jiao Tong University, Shanghai, China

{yuteng_ye, henryzh, jaile.cai, cxg, youjiazhang, yjqing, weiyangcs}@hust.edu.cn; wangjunle@gmail.com; qiang.hu@sjtu.edu.cn

Abstract

Occluded person re-identification (ReID) is a challenging problem due to contamination from occluders. Existing approaches address the issue with prior knowledge cues, such as human body key points and semantic segmentations, which easily fail in the presence of heavy occlusion and other humans as occluders. In this paper, we propose a feature pruning and consolidation (FPC) framework to circumvent explicit human structure parsing. The framework mainly consists of a sparse encoder, a multi-view feature matching module, and a feature consolidation decoder. Specifically, the sparse encoder drops less important image tokens, mostly related to background noise and occluders, solely based on correlation within the class token attention. Subsequently, the matching stage relies on the preserved tokens produced by the sparse encoder to identify k-nearest neighbors in the gallery by measuring the image and patch-level combined similarity. Finally, we use the feature consolidation module to compensate pruned features using identified neighbors for recovering essential information while disregarding disturbance from noise and occlusion. Experimental results demonstrate the effectiveness of our proposed framework on occluded, partial, and holistic Re-ID datasets. In particular, our method outperforms state-of-the-art results by at least 8.6% mAP and 6.0% Rank-1 accuracy on the challenging Occluded-Duke dataset.

Introduction

Person Re-Identification (ReID) refers to the process of retrieving the same person from a gallery set under non-overlapping surveillance cameras (Chen et al. 2017; Ye et al. 2021), and has been making remarkable progress in tackling appearance change in deep learning era (Wu et al. 2019; Lavi, Serj, and Ullah 2018). However, the re-identification of occluded persons remains a challenging problem because of two reasons: 1. the inference from wrongly included occluder features and 2. the partial or full absence of essential target features. To tackle occlusions, many existing approaches explicitly recover human semantics, via human pose estimation (Miao et al. 2019; Gao et al. 2020; Wang et al. 2022a) or body segmentation (Huang, Chen, and Huang 2020), as extra supervision to guide the network

to focus on non-occluded features. Others (Yu et al. 2021; Xu et al. 2022) first partition input image into horizontal or vertical parts, and then identify the occlusion status of each part with off-the-shelf models (Mask-RCNN (He et al. 2017), HR-Net (Sun et al. 2019)), and finally recover occluded features from K-nearest neighbors in a gallery according to non-occluded features. Both strategies rely on extra modules for detecting occlusions and will easily fail in the presence of heavy occlusion and persons as occluders, and further background noise persists as the partition usually is very coarse.

Inspired by the recent advances in sparse encoders (Rao et al. 2021; Liang et al. 2022), we propose a feature pruning, matching, and consolidation (FPC) framework for Occluded Person Re-Identification which adaptively removes interference from occluders and background and consolidates the contaminated features. Firstly, we send the query image into a modified transformer encoder with token sparsification to drop interference tokens (usually related to occluders and background) while preserving attentive tokens. Different from extra cue-based approaches that rely on prior information about human semantics, the sparse encoder exploits correlation properties on attention maps and generalizes better to various occlusion situations. In addition, our sparse encoder removes interference from the background as an extra benefit. Then, we rank the full tokens in the gallery memory according to their similarity with the query image. We obtain the gallery memory containing [cls] token and patch tokens via pre-training a vision transformer encoder. The similarity metric for matching is defined as the linear combination of image-level cosine distance and patch-level earth mover’s distance (Rubner, Tomasi, and Guibas 2000) for bridging the domain gap between the sparse query feature and holistic features in the gallery memory. At last, we select the k-nearest neighbors for each query and construct multi-view features by concatenating averaged [cls] tokens and respective patch tokens of both the query and its selected neighbors. We send the multi-view feature into a transformer decoder for feature consolidation. We compensate the occluded query feature with gallery neighbors and achieve better performance. During training, our method utilizes the entire training set as the gallery, a common practice in ReID literature (Xu et al. 2022; Yu et al. 2021). While for

*indicates corresponding author.

inference, only the test set query and gallery images are used without any pre-filtering steps. Our FPC framework achieves state-of-the-art performance on both occluded person, partial person, and holistic person ReID datasets. In particular, FPC outperforms the state-of-the-art by 8.6% mAP and 6.0% Rank-1 accuracy on the challenging Occluded-Duke dataset.

Our main contributions are as follows:

- We introduce the token sparsification mechanism to the occluded person ReID problem, the first to the best of our knowledge, which avoids explicit use of human semantics and better prunes unrelated features.
- We propose feature matching and consolidation modules to recover occluded query features from multi-view gallery neighbors. Compared to feature division approaches (Yu et al. 2021; Xu et al. 2022), our module uses transformer tokens that naturally preserve connectivity and richness of the features.
- We design a novel token based metric to measure the similarity of images by linearly combining the patch-level distance and the image-level distance.

Related Work

Our method is closely related to occluded person ReID and feature pruning in vision transformers.

Occluded Person ReID. The task of Occluded Person ReID is to find the same person under different cameras while the target pedestrian is obscured. The current methods are mainly divided into two categories, i.e., extra-clues based methods and feature reconstruction based methods. The extra-clues based methods locate non-occluded areas of the human body by prior knowledge cues, e.g., human pose (Miao et al. 2019; Gao et al. 2020; Wang et al. 2020a, 2022a; Miao, Wu, and Yang 2021) and human segments (Huang, Chen, and Huang 2020; Cai, Wang, and Cheng 2019). Another approach is based on feature reconstruction. Hou *et al.* (Hou et al. 2021) locate occluded human parts by key-points and propose Region Feature Completion (RFC) to recover the semantics of occluded regions in feature space. Xu *et al.* (Xu et al. 2022) extract occluded features with pose information and propose Feature Recovery Transformer (FRT) to recover the occluded features using the visible semantic features in the k-nearest gallery neighbors. In contrast, our approach adaptively removes interference from occluders and background according to correlation within the class token attention, which generalizes better to various occlusion situations.

Transformer Sparsification. The techniques of accelerating vision transformer models are necessary with limited computing resources. Most methods mainly aim to simplify the structure of the model with efficient attention mechanisms (Wang et al. 2020b; Kitaev, Kaiser, and Levskaya 2020; Zhang et al. 2022; Zhu et al. 2021) and compact structures (Liu et al. 2021; Touvron et al. 2021; Wu et al. 2021; Graham et al. 2021). There are also many researches (Kong et al. 2022; Rao et al. 2021; Liang et al. 2022) focus on effective token learning to reduce the redundancy. However, the above approaches have not explored the possibility of

applying the characteristics of model acceleration to the occluded person ReID problem. Actually, since there are various occlusion and background noise in occluded person ReID tasks, we observe that it is sub-optimal to apply the transformer model directly. Following (Liang et al. 2022), the purpose of our approach is to prune out ineffective tokens by exploring the sparsity of informative image patches.

Method

We illustrate our feature pruning and consolidation framework as illustrated in fig. 1. The framework consists of (1) sparse encoder \mathcal{S} conducts token sparsification to prune interference tokens and preserve attentive tokens; (2) multi-view feature matching module \mathcal{M} generates a rank list between the sparse query feature and pre-trained gallery memory by the image and patch-level combined similarity; (3) feature consolidation framework \mathcal{C} utilizes complete information of identified neighbors to compensate pruned query features.

Sparse Encoder

Inspired by the advances in feature pruning (Liang et al. 2022) and person ReID (He et al. 2021), we use a transformer with token sparsification to prune interference from occlusion and background. As shown in fig. 2, given an input image $x \in \mathbb{R}^{H \times W \times C}$, where W, H, C denote the width, height, and channel of the image respectively, we split x into N overlapping patches $\{p_1, p_2, \dots, p_N\}$ and embed each patch with linear projection denoted as $f(\cdot)$. Then we combined a learnable [cls] token with patch embedding and apply positional encoding and camera index encoding following (He et al. 2021). The final input can be described as:

$$\mathcal{Z} = \{x_{cls}, f(p_1), \dots, f(p_N)\} + \mathcal{P} + \mathcal{C}_{id} \quad (1)$$

where $x_{cls} \in \mathbb{R}^{1 \times D}$ is learnable [cls] token. $f(p_i) \in \mathbb{R}^{1 \times D}$ is i -th patch embeddings. $\mathcal{P} \in \mathbb{R}^{(N+1) \times D}$ is position embeddings and $\mathcal{C}_{id} \in \mathbb{R}^{(N+1) \times D}$ is camera index embeddings.

Token Sparsification. We adopt the token sparsification strategy proposed in (Liang et al. 2022). Specifically, through the attention correlation (Vaswani et al. 2017) between [cls] token and other tokens in the vision transformer, we can express the value of [cls] token as:

$$x_{cls} = A_{cls}V = \text{softmax}\left(\frac{Q_{cls}K}{\sqrt{d}}\right)V \quad (2)$$

where A_{cls} denotes the attention matrix of the [cls] token, i.e., first row of the attention matrix, \sqrt{d} is scale factor. Q_{cls}, K, V represent the query matrix of [cls] token, the key matrix, and the value matrix respectively. For multiple heads in the self-attention layer, we average the attention matrix as $\bar{A}_{cls} = \sum_i^n \frac{1}{n} \cdot A_{cls}^{(i)}$, n is the total heads number. Since the [cls] token corresponds to a larger attention value in significant patch regions (Caron et al. 2021), we can evaluate the importance of a token according to its relevance to the [cls] token. As \bar{A}_{cls} represents the correlation between [cls] token and all other tokens, we hence preserve the tokens with

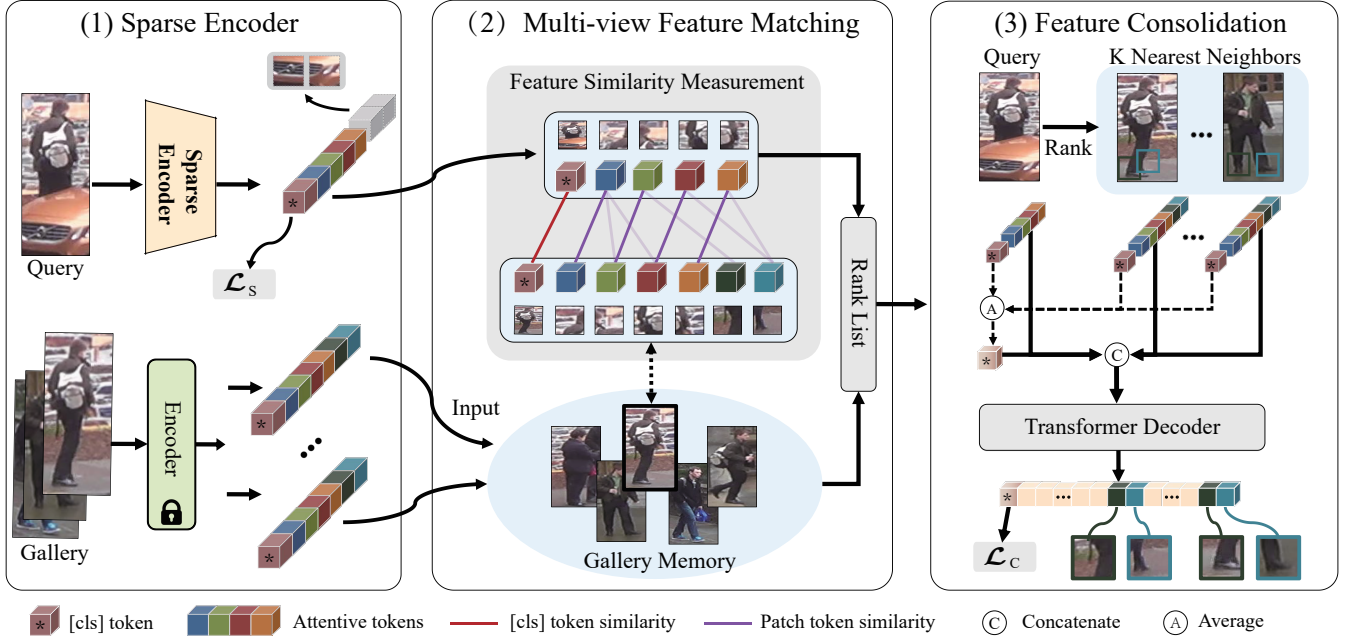


Figure 1: Overview of the proposed framework, which consists of a sparse encoder \mathcal{S} , a multi-view feature matching module \mathcal{M} , and a feature consolidation module \mathcal{C} . The sparse encoder \mathcal{S} removes interfering tokens while preserving attentive tokens. In matching module \mathcal{M} , we generate a rank list between the sparse query feature and holistic features in a gallery memory pretrained with a vision transformer. We use the summation of image-level cosine distance and patch-level earth mover’s distance as the metric for ranking. In the consolidation module \mathcal{C} , we select the k -nearest neighbors for the query and construct multi-view features by concatenating averaged [cls] tokens and respective patch tokens of both the query and its selected neighbors. The multi-view features are sent to the transformer decoder for feature consolidation.

\mathcal{K} largest values in A_{cls}^- and drop others, which is shown in fig. 2. We define $\mathcal{K} = \lceil \gamma \cdot N_c \rceil$, where γ is the keep rate and N_c is the total number of tokens in the current layer, $\lceil \cdot \rceil$ is the ceiling operation. With token sparsification, the preserved tokens are mostly related to the region of the target pedestrian, and dropped tokens are related to occluders or backgrounds.

Sparse Encoder Supervision Loss. We use the cross-entropy ID loss and triplet loss to supervise [cls] token x_{cls}^S obtained by the sparse encoder as:

$$\mathcal{L}_S = \mathcal{L}_{ID}(x_{cls}^S) + \mathcal{L}_T(x_{cls}^S) \quad (3)$$

where \mathcal{L}_{ID} denotes cross-entropy ID loss and \mathcal{L}_T denotes triplet loss. Compared with existing approaches, our feature pruning with token sparsification is adaptive and doesn’t rely on prior knowledge of human semantics, and can better handle challenging scenarios, such as heavy occlusion and other persons as occluders, as shown in fig. 6.

Multi-view Feature Matching Module

After the feature pruning, we would like to find the most related patches from other views that not been occluded for consolidation, such strategy has been proved to be effective (Xu et al. 2022; Yu et al. 2021). Specifically, we first learn a gallery memory with the pre-trained encoder. With the pruned query image features, we rank patches in the gallery memory according to their similarity with the query.

Considering that there exists appearance gap between the pruned query feature and holistic feature in gallery memory, we measure the similarity from both the image-level and patch-level. As shown in fig. 1, we linearly combine image-level and patch-level distance to match the query image with gallery memory images.

Image-level distance. We define the image-level distance as the cosine similarity between the [cls] tokens of query and gallery memory as follows:

$$\mathcal{D}_{COS} = 1 - \frac{\langle x_{cls}^{(i)}, x_{cls}^{(j)} \rangle}{|x_{cls}^{(i)}| \cdot |x_{cls}^{(j)}|} \quad (4)$$

where $x_{cls}^{(i)}$ and $x_{cls}^{(j)}$ is [cls] token of i -th query image and j -th image in gallery memory. $\langle \cdot \rangle$ is the dot product.

Patch-level distance. We leverage Earth Mover’s Distance (EMD) (Rubner, Tomasi, and Guibas 2000) to measure patch-level similarity. EMD is usually employed for measuring the similarity between two multidimensional distributions and is formulated as following linear programming problem: Let $\mathcal{Q} = \{(q_1, w_{q_1}), \dots, (q_m, w_{q_m})\}$ be the set of patch tokens from the query image, where q_i is i -th patch token and w_{q_i} is the weight of q_i . Similarly, $\mathcal{G} = \{(g_1, w_{g_1}), \dots, (g_n, w_{g_n})\}$ represents set of patch tokens of a gallery image. Then w_{q_i} is further defined as proportional correlation weight (Phan and Nguyen 2022) of q_i to set $\{g_1, g_2, \dots, g_n\}$, denoted as $\max(0, \langle q_i, \frac{1}{n} \sum_i g_i \rangle)$,

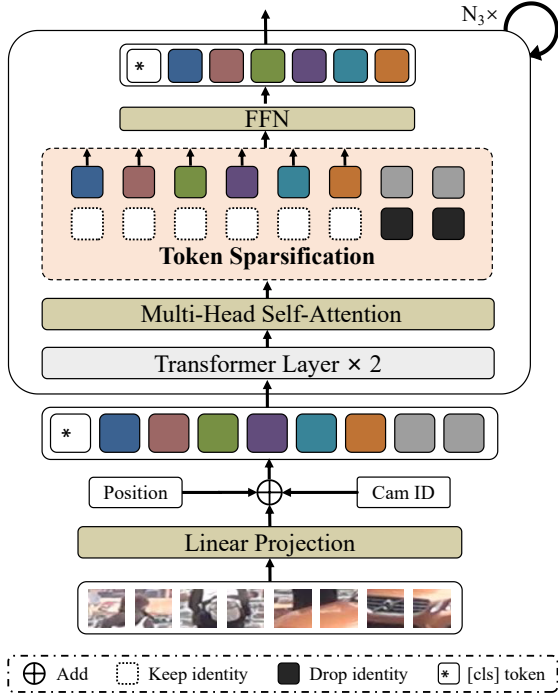


Figure 2: The structure of our sparse encoder. We divide and embed the image tokens using linear projection, positional encoding, and camera ID encoding. We perform token sparsification in layer 3, 6, and 9.

and w_{g_j} equals $\max(0, \langle g_j, \frac{1}{m} \sum_i q_i \rangle)$. The ground distance d_{ij} is the cosine distance between p_i and q_j . The objective is to find the flow \mathcal{F} , of which f_{ij} indicates the flow between p_i and q_j , to minimize following cost:

$$\begin{aligned} \mathcal{F}^* &= \arg \min_{\mathcal{F}} \text{COST}(\mathcal{Q}, \mathcal{G}, \mathcal{F}) \\ &= \arg \min_{\mathcal{F}} \left(\sum_{i=1}^m \sum_{j=1}^n f_{ij} d_{ij} \right) \end{aligned} \quad (5)$$

subject to the following constraints:

$$f_{ij} \geq 0, \quad 1 \leq i \leq m, \quad 1 \leq j \leq n \quad (6)$$

$$\sum_{i=1}^m f_{ij} \leq w_{g_j}, \quad 1 \leq j \leq n \quad (7)$$

$$\sum_{j=1}^n f_{ij} \leq w_{q_i}, \quad 1 \leq i \leq m \quad (8)$$

$$\sum_{i=1}^m \sum_{j=1}^n f_{ij} = \min \left(\sum_{i=1}^m w_{q_i}, \sum_{j=1}^n w_{g_j} \right) \quad (9)$$

eq. (5) can be solved by iterative Sinkhorn algorithm (Curi 2013) and produces the earth mover's distance \mathcal{D}_{EMD} . Intuitively, \mathcal{D}_{COS} represents global distance, while \mathcal{D}_{EMD} measures in the perspective of the set of local features. Naturally, we take the linear combination of image-level cosine

distance and patch-level EMD as our final distance for feature matching, expressed as follows:

$$\mathcal{D} = (1 - \alpha) \mathcal{D}_{COS} + \alpha \mathcal{D}_{EMD} \quad (10)$$

With \mathcal{D} , we rank the sparse query feature with holistic features in gallery memory and generate a ranking list. In order to expedite the process of feature matching, we use an efficient two-stage selection strategy to save time costs by 138 times.

Feature Consolidation Decoder

After we retrieve the K nearest gallery neighbors, we then consolidate the pruned query feature from multi-view observations in the gallery memory. Specifically, we average the [cls] tokens of the query and gallery neighbors as our initial global feature. Then we combine the averaged [cls] token with patch tokens in both query and gallery neighbors to aggregate information of multi-view pedestrians. The consolidation of the multi-view features can be formulated as follows:

$$f_m = \left\{ \bar{f}_c, f_{p_q}, f_{p_g}^{(1)}, f_{p_g}^{(2)}, \dots, f_{p_g}^{(K)} \right\} \quad (11)$$

where $\bar{f}_c \in \mathbb{R}^{1 \times C}$ is the average of [cls] tokens of both query and its neighbors. $f_{p_q} \in \mathbb{R}^{M \times C}$ is the patch tokens of query and M is the number of patch tokens. $f_{p_g}^{(i)} \in \mathbb{R}^{N \times C}$ is the patch tokens of i -th gallery neighbor and N is the number of patch tokens for each gallery.

Transformer Decoder. Notice that [cls] token is class prediction (Dosovitskiy et al. 2020) in transformer and conventionally treated as global feature for image representation, the query [cls] token is usually contaminated in the occluded person ReID problem. We send the consolidated multi-view feature to a transformer decoder to compensate the incomplete [cls] token with gallery neighbors. In the decoder, we first transform the multi-view feature f_m into Q, K, V vectors using linear projections, which is:

$$Q = W_q \cdot f_m, K = W_k \cdot f_m, V = W_v \cdot f_m \quad (12)$$

where W_q, W_k, W_v are weights of linear projections. As the multi-view feature contains three parts: [cls] token, patch tokens of query, and patch tokens of gallery neighbors, the computation of attention with respect to [cls] token can be decomposed into the above three parts, which is expressed as:

$$\begin{aligned} x_{cls} &= \text{Cat}(A'_c, A'_q, A'_g) \cdot \text{Cat}(V_c, V_q, V_g) \\ &= A'_c \cdot V_c + A'_q \cdot V_q + A'_g \cdot V_g \end{aligned} \quad (13)$$

where $\text{Cat}(\cdot)$ is operation of vector combination. A' denotes the attention matrix of [cls] token and V is the value vector. Subscripts c, q, g correspond to [cls] token, query and gallery neighbors respectively. $A'_c \cdot V_c + A'_q \cdot V_q$ acts as the feature learning process with sparse query and $A'_g \cdot V_g$ integrates completion information from gallery neighbors to the [cls] token. The final consolidated [cls] token generated by the transformer decoder is denoted as $x_{cls}^c = \tau(f_m)$. We find one layer of transformer is enough and also avoids the

high memory consumption from neighborhoods (Yu et al. 2021).

Consolidation Loss. The [cls] token x_{cls}^c obtained by feature consolidation decoder is supervised with cross entropy loss as ID loss and triplet loss, as expressed below:

$$\mathcal{L}_C = \mathcal{L}_{ID}(x_{cls}^c) + \mathcal{L}_T(x_{cls}^c) \quad (14)$$

Therefore, the final loss function can be expressed:

$$\mathcal{L} = \mathcal{L}_S + \mathcal{L}_C \quad (15)$$

Implementation Details

We choose the ViT-B/16 as our backbone for both the sparse encoder and gallery encoder. Our sparse encoder incorporates several modifications into the backbone, including camera encoding, patch overlapping with a stride of 11, batch normalization, and token sparsification at layer 3, 6, 9. Our gallery encoder has the same structure as the sparse encoder but does not conduct token sparsification. We construct the gallery memory with gallery image tokens as in eq. (1). In the multi-view feature matching module, we first identify 100 globally-similar gallery neighbors using cos distance in eq. (4), and then search the final K nearest gallery neighbors using the proposed distance in eq. (10). We set the number of K to 10 on Occluded-Duke dataset and 5 for others. In the training process, we resize all input images to 256×128 . The training images are augmented with random horizontal flipping, padding, random cropping and random erasing (Zhong et al. 2020). The batch size is 64 with 4 images per ID and the learning rate is initialized as 0.008 with cosine learning rate decay. The distance weight α in eq. (10) is 0.4 and the parameter of keep rate γ is 0.8. We take consolidated [cls] token in eq. (13) for model inference. For large-scale problems, we can further replace the full image tokens with [cls] tokens to save computation and memory.

Experiment

We conduct extensive experiments to validate our framework. We first introduce the dataset we use:

Datasets and Evaluation Metric

Occluded-Duke (Miao et al. 2019) includes 15,618 training images of 702 persons, 2,210 occluded query images of 519 persons, and 17,661 gallery images of 1,110 persons. **Occluded-ReID** (Zhuo et al. 2018) consists of 1,000 occluded query images and 1,000 full-body gallery images both belonging to 200 identities. **Partial-ReID** (Zheng et al. 2015b) involves 600 images from 60 persons, and each person consists 5 partial and 5 full-body images. **Market-1501** (Zheng et al. 2015a) contains 12,936 training images of 751 persons, 19,732 query images and 3,368 gallery images of 750 persons captured by six cameras.

Evaluation Metric. All methods are evaluated under the Cumulative Matching Characteristic (CMC) and mean Average Precision (mAP). Floating Point Operations (FLOPs) represents the amount of model computation.

Method	Occluded-Duke		Occluded-ReID	
	Rank-1	mAP	Rank-1	mAP
DSR	40.8	30.4	72.8	62.8
PGFA	51.4	37.3	-	-
HOReID	55.1	43.8	55.1	43.8
OAMN	62.6	46.1	-	-
PAT	64.5	53.6	81.6	72.1
TransReID	67.4	59.5	-	-
PFD	69.5	61.8	81.5	83.0
FED	67.9	56.3	87.0	79.4
RFCnet	63.9	54.5	-	-
Yu et al.	67.6	64.2	68.8	67.3
FRT	70.7	61.3	80.4	71.0
FPC (ours)	76.7	72.8	86.3	84.6

Table 1: Comparison with state-of-the-art methods on Occluded-Duke and Occluded-ReID datasets.

Method	Market-1501		Partial-ReID	
	Rank-1	mAP	Rank-1	mAP
PCB	92.3	77.4	-	-
PGFA	91.2	76.8	69.0	61.5
HOReID	94.2	84.9	85.3	-
OAMN	92.3	79.8	86.0	-
PAT	95.4	88.0	88.0	-
TransReID	95.0	88.2	83.0	77.5
PFD	95.5	89.7	-	-
FED	95.0	86.3	84.6	82.3
RFCnet	95.2	89.2	-	-
Yu et al.	94.5	86.5	-	-
FRT	95.5	88.1	88.2	-
FPC (ours)	95.1	91.4	86.3	86.5

Table 2: Comparison with state-of-the-art methods on Market-1501 and Partial-ReID datasets.

Comparison with State-of-the-art Methods

Experimental results on Occluded ReID Datasets. In table 1, we compare FPC with state-of-the-art methods on two occluded ReID datasets (i.e., Occluded-Duke and Occluded-ReID). Methods in different categories are compared, including the partial ReID methods (He et al. 2018), key-points based methods (Miao et al. 2019; Wang et al. 2020a, 2022a; Hou et al. 2021), data augmentation methods (Chen et al. 2021; Wang et al. 2020a, 2022b), transformer-based methods (Li et al. 2021; He et al. 2021), gallery-based reconstruction methods (Yu et al. 2021; Xu et al. 2022). In addition, TransReID (He et al. 2021) and PFD (Wang et al. 2022a) use camera information. Since no training set is provided for Occluded-ReID, we adopt Market-1501 as the training set the same as other methods to ensure the fairness of comparison. We can see that our FPC outperforms existing approaches and demonstrates the effectiveness for the occluded ReID tasks. Specifically, our FPC achieves the best performance on the challenging Occluded-Duke dataset, outperforming other methods by at least 6.0% Rank-1 accuracy and 8.6% mAP. On the Occluded-ReID dataset,

Index	\mathcal{S}	\mathcal{M}	\mathcal{C}	Rank-1	mAP
1	✗	✗	✗	65.2	56.6
2	✓	✗	✗	68.6	60.1
3	✓	✗	✓	75.8	71.6
4	✓	✓	✓	76.7	72.8

Table 3: Ablation study on each component.

FPC produces the highest mAP, outperforming the other methods by at least 1.6%. FPC achieves comparable results in Rank-1 accuracy with FED (Wang et al. 2022b), and much better than others.

Experimental Results on Holistic and Partial ReID Datasets. Many existing occluded ReID methods can not be effectively applicable to holistic and partial ReID datasets. On the contrary, FPC achieves great performance improvement on holistic and partial datasets, i.e., Market-1501 and Partial-ReID. The results are shown in table 2. We compare FPC with three categories of methods, including the holistic ReID methods (Sun et al. 2018), the current leading methods in occluded ReID (Miao et al. 2019; Wang et al. 2020a; Chen et al. 2021; Li et al. 2021; He et al. 2021; Wang et al. 2022a,b), the feature reconstruction based methods in occluded ReID (Hou et al. 2021; Yu et al. 2021; Xu et al. 2022). On the Partial-ReID dataset, FPC achieves the best mAP result, lower than the PAT and FRT in Rank-1 accuracy. We consider that we adopt the ViT models and are more prone to overfit on small datasets which leads poor cross-domain generalization. We also observe that FPC achieves competitive Rank-1 accuracy and the highest mAP on the Market-1501 dataset, which is at least 1.7% mAP ahead of other methods. The excellent performance on holistic and partial ReID datasets illustrates the robustness of our FPC.

Ablation Study

Analysis of proposed components. The results are shown in table 3. Index-1 is our baseline architecture. To assess the effectiveness of \mathcal{S} , we conduct a comparison between index-1 and index-2. Our findings demonstrate that \mathcal{S} leads to a 3.4% improvement in Rank-1 accuracy and a 3.5% increase in mAP over the baseline. This suggests that \mathcal{S} has the potential to mitigate the interference of inattentive features (mainly related to occlusion and background noise). Furthermore, \mathcal{S} can also expedite model inference, as elaborated in table 4. By comparing index-2 and index-3, \mathcal{C} improves performance by 7.2% Rank-1 accuracy and 11.5% mAP, which indicates that \mathcal{C} can effectively compensate the occluded query feature and bring huge performance improvements. By comparing index-3 and index-4, \mathcal{M} can increase performance by 0.9% Rank-1 accuracy and 1.2% mAP, indicating that the proposed distance metric contributes to the find the accurate gallery neighbors.

Analysis of Distance Weight α . The distance weight α defined in eq. (10) balances the importance between image-level and patch-level distance. From fig. 3a, as α goes from 0 to 1, the patch-level distance gradually takes effect. When α is 0.4, the combination of both achieves the best performance with 76.7% Rank-1 accuracy and 72.8% mAP.



Figure 3: Analysis of distance weight α and the number of nearest neighbors K on Occluded-Duke dataset. mAP and Rank-1 are respectively depicted in blue and orange colors.

Figure 4: Illustration of three different patch drop strategies. White patches are the dropped image parts.

Analysis of Number K of Nearest Gallery Neighbors.

We conduct quantitative experiments to find the most appropriate number K of nearest gallery neighbors. We conclude from fig. 3b that too small a choice of K lacks sufficient information for feature consolidation and too large a choice increases the risk of incorrectly selecting neighbors. When the value of K is 10, we achieve the optimal performance.

Analysis of Keep Rate γ in Sparse Encoder. Here, keep rate γ reflects the number of preserved tokens in each layer of the sparse encoder. We find instructive observations from table 4: as the γ goes from 1.0 to 0.8, the FLOPs drops while the model performance improves. This suggests that a reasonable choice of γ can effectively filter out inattentive features, thus reducing the computational complexity and enhancing the model inference capability. When γ is 0.8, we achieve the best balance between computational complexity and experimental performance, which leads to an improvement of 0.5% mAP, 0.7% Rank-1 accuracy and 25% reduction in FLOPs.

Analysis of Patch Drop Strategy in Sparse Encoder.

We experiment with three variants of patch drop approaches to demonstrate the effectiveness of the proposed method. As shown in fig. 4, we preserve the same number of patches and choose different preservation methods: Random Drop means that K patches are randomly selected to be retained. Non-salient Drop, as the proposed method, preserves the patches corresponding to the K largest values in the [cls] token attention. Conversely, Salient Drop preserves the K smallest ones. The performance of the three patch drop approaches with different keep rate is shown in fig. 5. We observe that the proposed Non-salient Drop method achieves best performance, indicating the importance of attentive features for

Metrics	Keep Rate γ					
	1.0	0.9	0.8	0.7	0.6	0.5
Rank-1 (%)	76.0	76.6 (+0.6)	76.7 (+0.7)	75.9 (-0.1)	74.5 (-1.5)	72.1 (-3.9)
mAP (%)	72.3	72.7 (+0.4)	72.8 (+0.5)	72.4 (+0.1)	71.3 (-1.0)	69.3 (-3.0)
FLOPs (G)	20.8	18.1 (-13%)	15.7 (-25%)	13.6 (-33%)	11.9 (-43%)	10.4 (-50%)


Table 4: Analysis on effectiveness of the keep rate γ . We perform FLOPs comparison on the sparse encoder.

Figure 5: The analysis of three different patch drop strategies. Blue, red, and orange colors represent the Non-salient Drop, Random Drop, and Salient Drop, respectively.

feature matching and consolidation. In addition, the result of random drop test achieves relatively high performance, which reflects our proposed sparse encoder structure is robust against information loss.

Analysis of speed and storage optimization. Our method relies on gallery information for feature consolidation, to address the practical concerns regarding speed and storage for real-world problems, we conduct analysis on our optimization strategies on the Occluded-Duke dataset. In \mathcal{S} , token sparsification reduce 25% Flops (in table 4) with γ equals 0.8 and thus accelerates model inference. In \mathcal{M} , we use a two-step matching procedure, i.e., use cosine similarity for initial filtering and compute EMD in the finer step, achieving a remarkable 138-fold reduction in computational costs, compared to the exhaustive search based on EMD. Our approach takes 0.43ms per image, while the full EMD calculation takes 60.0ms per image. Furthermore, if we replace full image tokens with $[cls]$ token, our approach saves a lot memory (i.e., only takes 0.05G for the entire gallery set) and achieves 17x faster in nearest neighbor search time (i.e., from 0.43ms to 0.026ms per image), while still preserving an acceptable degree of precision (i.e., mAP: 72.8% to 70.2%, Rank-1: 76.7% to 74.3%, which still outperforms others).

Comparisons with Re-ranking. An alternative approach is the re-ranking technique (Zhong et al. 2017), which also uses k-nearest gallery neighbors information. We compare \mathcal{C} with re-ranking (Zhong et al. 2017), the second group results in table 5 indicate that re-ranking fails to attain comparable performance of \mathcal{C} , lead to a reduction of 3.2% Rank-1 and 0.3% mAP. Further, our $\mathcal{M} + \mathcal{C}$ is approximately 17 times faster than re-ranking, with respective times of 0.47ms and 7.9ms per image. Moreover, the last group shows \mathcal{C} and re-ranking can be jointly employed and outperforms others.

Methods	Rank-1	mAP
FRT(Xu et al. 2022)	70.7	61.3
FRT(Xu et al. 2022) + re-ranking	70.8	65.0
Yu et al. (Yu et al. 2021)	67.6	64.2
Yu et al. (Yu et al. 2021) + re-ranking	68.9	67.3
\mathcal{S}	68.6	60.1
\mathcal{S} + re-ranking	72.6	71.3
$\mathcal{S} + \mathcal{C}$	75.8	71.6
FPC (ours)	76.7	72.8
FPC (ours) + re-ranking	78.6	78.3


Table 5: Performance comparisons with re-ranking on Occluded-Duke dataset. \mathcal{S} is our sparse encoder and \mathcal{C} is feature consolidation.

Figure 6: Visualization of patch drop process in different layers of the sparse encoder. We show various occlusion scenarios such as object occlusion, pedestrian occlusion and heavy occlusions.

Visualization

Patch pruning. The patch-drop process in different layers of sparse encode is shown in fig. 6. We observe that with the increasing number of sparse encoder layers, more object occlusion, pedestrian occlusion and background noise are filtered out, while the essential classification information of target pedestrians is preserved.

Conclusion

In this paper, we propose the feature pruning and consolidation (FPC) framework for the occluded person ReID task. Specifically, our framework prunes interfering image tokens (mostly related to background noise and occluders) without relying on prior human shape information. We propose an effective way to consolidate the pruned feature and achieve the SOTA performance on many occluded, partial, and holistic ReID datasets. We introduce the token sparsification technique and demonstrate its effectiveness.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (NSFC No. 62272184). The computation is completed in the HPC Platform of Huazhong University of Science and Technology.

References

- Cai, H.; Wang, Z.; and Cheng, J. 2019. Multi-scale body-part mask guided attention for person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 0–0.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9650–9660.
- Chen, P.; Liu, W.; Dai, P.; Liu, J.; Ye, Q.; Xu, M.; Chen, Q.; and Ji, R. 2021. Occlude them all: Occlusion-aware attention network for occluded person re-id. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11833–11842.
- Chen, Y.-C.; Zhu, X.; Zheng, W.-S.; and Lai, J.-H. 2017. Person re-identification by camera correlation aware feature augmentation. *IEEE transactions on pattern analysis and machine intelligence*, 40(2): 392–408.
- Cuturi, M. 2013. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Gao, S.; Wang, J.; Lu, H.; and Liu, Z. 2020. Pose-guided visible part matching for occluded person reid. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11744–11752.
- Graham, B.; El-Nouby, A.; Touvron, H.; Stock, P.; Joulin, A.; Jégou, H.; and Douze, M. 2021. LeViT: a Vision Transformer in ConvNet’s Clothing for Faster Inference. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 12239–12249.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, 2961–2969.
- He, L.; Liang, J.; Li, H.; and Sun, Z. 2018. Deep spatial feature reconstruction for partial person re-identification: Alignment-free approach. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7073–7082.
- He, S.; Luo, H.; Wang, P.; Wang, F.; Li, H.; and Jiang, W. 2021. Transreid: Transformer-based object re-identification. In *Proceedings of the IEEE/CVF international conference on computer vision*, 15013–15022.
- Hou, R.; Ma, B.; Chang, H.; Gu, X.; Shan, S.; and Chen, X. 2021. Feature completion for occluded person re-identification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Huang, H.; Chen, X.; and Huang, K. 2020. Human parsing based alignment with multi-task learning for occluded person re-identification. In *2020 IEEE International Conference on Multimedia and Expo (ICME)*, 1–6. IEEE.
- Kitaev, N.; Kaiser, L.; and Levskaya, A. 2020. Reformer: The Efficient Transformer. In *8th International Conference on Learning Representations (ICLR)*.
- Kong, Z.; Dong, P.; Ma, X.; Meng, X.; Niu, W.; Sun, M.; Shen, X.; Yuan, G.; Ren, B.; Tang, H.; Qin, M.; and Wang, Y. 2022. SPViT: Enabling Faster Vision Transformers via Latency-Aware Soft Token Pruning. In *Proceedings of the European conference on computer vision (ECCV)*.
- Lavi, B.; Serj, M. F.; and Ullah, I. 2018. Survey on deep learning techniques for person re-identification task. *arXiv preprint arXiv:1807.05284*.
- Li, Y.; He, J.; Zhang, T.; Liu, X.; Zhang, Y.; and Wu, F. 2021. Diverse part discovery: Occluded person re-identification with part-aware transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2898–2907.
- Liang, Y.; Ge, C.; Tong, Z.; Song, Y.; Wang, J.; and Xie, P. 2022. Not All Patches are What You Need: Expediting Vision Transformers via Token Reorganizations. In *arXiv preprint arXiv:2202.07800*.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 10012–10022.
- Miao, J.; Wu, Y.; Liu, P.; Ding, Y.; and Yang, Y. 2019. Pose-guided feature alignment for occluded person re-identification. In *Proceedings of the IEEE/CVF international conference on computer vision*, 542–551.
- Miao, J.; Wu, Y.; and Yang, Y. 2021. Identifying visible parts via pose estimation for occluded person re-identification. *IEEE Transactions on Neural Networks and Learning Systems*.
- Phan, H.; and Nguyen, A. 2022. DeepFace-EMD: Re-Ranking Using Patch-Wise Earth Mover’s Distance Improves Out-of-Distribution Face Identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20259–20269.
- Rao, Y.; Zhao, Y.; Liu, B.; Lu, J.; Zhou, J.; and Hsieh, C.-J. 2021. Dynamicvit: Efficient vision transformers with dynamic token sparsification. In *Advances in Neural Information Processing Systems (NeurIPS)*, 13937–13949.
- Rubner, Y.; Tomasi, C.; and Guibas, L. J. 2000. The earth mover’s distance as a metric for image retrieval. *International journal of computer vision*, 40(2): 99–121.
- Sun, K.; Xiao, B.; Liu, D.; and Wang, J. 2019. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5693–5703.

- Sun, Y.; Zheng, L.; Yang, Y.; Tian, Q.; and Wang, S. 2018. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In *Proceedings of the European conference on computer vision (ECCV)*, 480–496.
- Touvron, H.; Cord, M.; Douze, M.; Massa, F.; Sablayrolles, A.; and Jegou, J. 2021. Training data-efficient image transformers & distillation through attention. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, 10347–10357.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, G.; Yang, S.; Liu, H.; Wang, Z.; Yang, Y.; Wang, S.; Yu, G.; Zhou, E.; and Sun, J. 2020a. High-order information matters: Learning relation and topology for occluded person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6449–6458.
- Wang, S.; Li, B.; Khabsa, M.; Fang, H.; and Ma, H. 2020b. Linformer: Self-Attention with Linear Complexity. In *arXiv preprint arXiv:2006.04768*.
- Wang, T.; Liu, H.; Song, P.; Guo, T.; and Shi, W. 2022a. Pose-guided feature disentangling for occluded person re-identification based on transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 2540–2549.
- Wang, Z.; Zhu, F.; Tang, S.; Zhao, R.; He, L.; and Song, J. 2022b. Feature Erasing and Diffusion Network for Occluded Person Re-Identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4754–4763.
- Wu, D.; Zheng, S.-J.; Zhang, X.-P.; Yuan, C.-A.; Cheng, F.; Zhao, Y.; Lin, Y.-J.; Zhao, Z.-Q.; Jiang, Y.-L.; and Huang, D.-S. 2019. Deep learning-based methods for person re-identification: A comprehensive review. *Neurocomputing*, 337: 354–371.
- Wu, H.; Xiao, B.; Codella, N.; Liu, M.; Xiyang, D.; Yuan, L.; and Zhang, L. 2021. CvT: Introducing Convolutions to Vision Transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 22–31.
- Xu, B.; He, L.; Liang, J.; and Sun, Z. 2022. Learning Feature Recovery Transformer for Occluded Person Re-Identification. *IEEE Transactions on Image Processing*, 31: 4651–4662.
- Ye, M.; Shen, J.; Lin, G.; Xiang, T.; Shao, L.; and Hoi, S. C. 2021. Deep learning for person re-identification: A survey and outlook. *IEEE transactions on pattern analysis and machine intelligence*, 44(6): 2872–2893.
- Yu, S.; Chen, D.; Zhao, R.; Chen, H.; and Qiao, Y. 2021. Neighbourhood-guided feature reconstruction for occluded person re-identification. *arXiv preprint arXiv:2105.07345*.
- Zhang, J.; Peng, H.; Wu, K.; Liu, M.; Xiao, B.; Fu, J.; and Yuan, L. 2022. MiniViT: Compressing Vision Transformers with Weight Multiplexing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12145–12154.
- Zheng, L.; Shen, L.; Tian, L.; Wang, S.; Wang, J.; and Tian, Q. 2015a. Scalable person re-identification: A benchmark. In *Proceedings of the IEEE international conference on computer vision*, 1116–1124.
- Zheng, W.-S.; Li, X.; Xiang, T.; Liao, S.; Lai, J.; and Gong, S. 2015b. Partial person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, 4678–4686.
- Zhong, Z.; Zheng, L.; Cao, D.; and Li, S. 2017. Re-ranking person re-identification with k-reciprocal encoding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1318–1327.
- Zhong, Z.; Zheng, L.; Kang, G.; Li, S.; and Yang, Y. 2020. Random erasing data augmentation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 13001–13008.
- Zhu, M.; Han, K.; Tang, Y.; and Wang, Y. 2021. Visual Transformer Pruning. In *arXiv preprint arXiv:2104.08500*.
- Zhuo, J.; Chen, Z.; Lai, J.; and Wang, G. 2018. Occluded person re-identification. In *2018 IEEE International Conference on Multimedia and Expo (ICME)*, 1–6. IEEE.