

Dense Interaction Learning for Video-based Person Re-identification

Tianyu He¹, Xin Jin², Xu Shen¹, Jianqiang Huang¹, Zhibo Chen², and Xian-Sheng Hua¹

¹DAMO Academy, Alibaba Group

²University of Science and Technology of China

timhe.hty@alibaba-inc.com

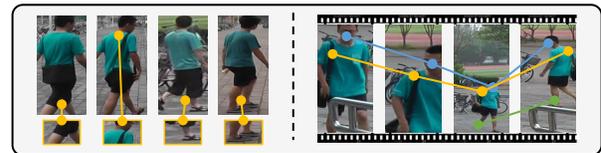
Abstract

Video-based person re-identification (*re-ID*) aims at matching the same person across video clips. Efficiently exploiting multi-scale fine-grained features while building the structural interaction among them is pivotal for its success. In this paper, we propose a hybrid framework, Dense Interaction Learning (DenseIL), that takes the principal advantages of both CNN-based and Attention-based architectures to tackle video-based person re-ID difficulties. DenseIL contains a CNN encoder and a Dense Interaction (DI) decoder. The CNN encoder is responsible for efficiently extracting discriminative spatial features while the DI decoder is designed to densely model spatial-temporal inherent interaction across frames. Different from previous works, we additionally let the DI decoder densely attends to intermediate fine-grained CNN features and that naturally yields multi-grained spatial-temporal representation for each video clip. Moreover, we introduce Spatio-TEmporal Positional Embedding (STEP-Emb) into the DI decoder to investigate the positional relation among the spatial-temporal inputs. Our experiments consistently and significantly outperform all the state-of-the-art methods on multiple standard video-based person re-ID datasets.

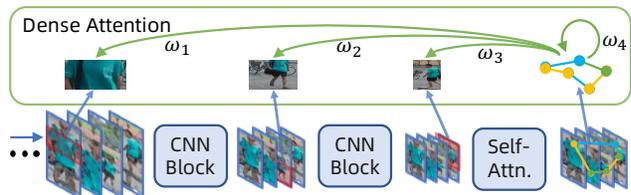
1. Introduction

Person re-identification (*re-ID*) tells whether a person-of-interest has been noticed in a different location by another camera. It is essential to many important surveillance applications such as tracking [70] and retrieval [89]. In recent years, significant progresses have been achieved in image-based person re-ID [42, 59, 63, 22], as well as the video-based one [87, 41, 80], due to the rapid development of Convolutional Neural Networks (CNN) [36].

The goal of image-based person re-ID is to match person images captured in different times and locations. To achieve this goal, recent methods are proposed to better dig appearance features while concurrently to maintain the robustness with respect to body part misalign-



(a) Left: different identities with similar appearance. Right: the same identity with misalignment or occlusion.



(b) Our proposed Dense Interaction Learning (DenseIL).

Figure 1: (a): Two key challenges existed in video-based person re-ID. (b): To tackle these two challenges, we introduce a DenseIL framework to densely capture multi-grained spatial-temporal interaction with the guidance of global relationship, where the preferences (ω) are automatically learned by our proposed Dense Attention.

ment [43, 54, 7, 91, 73, 63, 67, 33, 32, 22]. Nevertheless, it is still insufficient to model the discrete relationships among various misaligned body parts from a single image. This motivates the exploration on video-based re-ID, which further takes adjacent frames of the captured person into consideration, and therefore the occluded or missed parts can be inferred. In this sense, the key point to video-based re-ID lies in designing an architecture that is suitable for temporal dynamics, such as leveraging optical flow [9], RNN [92] and 3D CNN [40].

However, the aforementioned methods neglect the importance of spatial-temporal interaction between body parts within intra- and inter-frame, which limits their effectivity. As a result, the state-of-the-arts [80, 79] suggest that graph convolutional network [34] has the merit of modeling spatial-temporal dependencies and shows promising performance on video-based re-ID task. However, their models are built upon the *coarse-grained* representation while leaving the

fine-grained information implied in each frame not fully exploited. As demonstrated in Figure 1a, when different identities share similar appearance, depending on coarse-grained knowledge is not enough to distinguish the difference. Instead, fine-grained information (such as a shoe, a shoulder bag, *etc.*) plays an enormous role in re-identification.

Inspired by this observation, in this paper, we present Dense Interaction Learning (DenseIL) that not only builds *spatial-temporal interaction* between body parts but also densely exploits *fine-grained cues* (see Figure 1b). Basically, DenseIL is composed of a CNN encoder and a Dense Interaction (DI) decoder. The CNN encoder exerts its advantage on efficiently encoding spatial context into discriminative features [37]. Our CNN encoder consists of several CNN Blocks (*e.g.*, Res-Block [21], Dense-Block [31], SE-Block [30], *etc.*) and therefore is capable of generating a set of hidden features from low-level/high-resolution to high-level/low-resolution.

Our DI decoder comprises stacked self-attention [66], feed-forward layer, layer normalization [1] and a newly proposed Dense Attention. Specifically, our decoder reuses the self-attention module explored in vanilla Transformer to deliberately model spatial-temporal inherent interaction across frames. After that, the subsequent Dense Attention module simultaneously attends to *both* the outputs of the self-attention module and the intermediate fine-grained CNN features of the CNN encoder with the guidance of global relationship (*i.e.*, the last self-attention outputs). Thus, the decoder can naturally generate a multi-grained spatial-temporal representation for each video clip.

In contrast to ResNet [21] and DenseNet [31] that associate features of preceding layers through *summation* or *concatenation*, we create dense information flow between CNN and Attention mechanism with the proposed Dense Attention scheme. Intuitively, our Dense Attention is designed to better facilitate the coordination of CNN and Attention mechanism by extracting hybrid information from both the convolution and preceding self-attention features in a single attention function. Acting in this way, the model automatically and flexibly learns the preference on the fine-grained context (frame level, extracted by CNN) and global spatial-temporal interaction across frames (sequence level, modeled by self-attention) when re-identifying person-of-interest, which will be beneficial to scale-variations and misalignment problem, as a by-product of our model design.

Moreover, since the DI decoder is intrinsically permutation-invariant, we further propose Spatio-TEmporal Positional Embedding (STEP-Emb) to explicitly enhance the chronological relation in intra- and inter-frame. By incorporating STEP-Emb, our DI decoder is able to investigate the absolute or relative position of the spatial-temporal inputs.

Experiments show that our method significantly outperforms the state-of-the-arts on three video-based person

re-ID datasets. In particular, we achieve the best performance of 87.0% and 97.1% mAP on large-scale MARS and DukeMTMC-VideoReID datasets respectively.

2. Background

2.1. Related Work

Video-based Person Re-identification. Video-based person re-identification (re-ID) [87] targets on re-identifying a video clip of the person-of-interest from massive gallery candidates. The majority of the video-based re-ID algorithms pay close attention to efficiently modeling temporal context and exploiting complementary information from various frames, by leveraging optical flow [51, 11, 9], RNN [78, 92, 77, 47], pooling [87, 74], 3D CNN [55, 40, 20], attention mechanism [41, 17, 45, 39, 85] and graph convolutional network [80, 79]. For example, McLaughlin *et al.* [51] utilize optical flow, RNN and temporal pooling simultaneously to take long and short term temporal cues into account. Qiu *et al.* [55] and Li *et al.* [40] both employ 3D CNN to attain spatial-temporal representation for each video clip. However, both optical flow and 3D CNN are sensitive to spatial misalignment between adjacent video frames.

More recently, the rise of attention mechanism [2] convincingly demonstrates high capability of selectively focusing on specific parts of the input signal. Inspired by this, lots of studies learn the weight of spatial and temporal features separately from a static perspective [46, 92, 77, 41, 17, 27, 39, 61, 26]. However, they do not take fully advantage of the relationships between spatial and temporal body parts across different frames, thus yielding limited capability of representation.

Attention Mechanism in Computer Vision Community.

The attention mechanism is first proven to be helpful in Neural Machine Translation [2, 66], of which Transformer [66] is the most famous one. Existing works in computer vision community leverage the success of Transformer mainly by replacing the convolutions with self-attention operation [53, 68], or by deploying attention as an add-on to existing convolutional models in tasks like image classification [30, 29, 3, 28], object detection [71, 4], segmentation [16, 83], generative networks [84], inpainting [82], *etc.*

The tale of attention mechanism still continues. People start to consider borrowing the entire Transformer architecture to jointly modeling vision-language representations [62, 49, 60, 56, 12] or exploiting relations of the objects in image object detection [5, 94]. While different from aforementioned studies that build the entire Transformer *on the highest-level of* CNN spatial features, we only engage the decoder of Transformer, and replace the vanilla encoder-decoder attention with the proposed Dense Attention to pay attention to multi-grained CNN representations.

Encoder-Decoder Framework with Skip Connections.

The encoder-decoder framework is widely applied in the areas of language [10, 2], speech [6] and vision [48, 93]. For image and video processing, Ronneberger *et al.* [58] first introduce concatenation-style skip connections [21, 31] to the encoder-decoder-based fully convolutional networks [48], and such framework is further utilized and refined in various computer vision tasks [50, 15, 64].

In this work, instead of using the conventional summation-style or concatenation-style skip connections like above methods, we introduce Dense Attention to *densely attend* to multi-grained features generated by the CNN encoder or the preceding decoder blocks.

2.2. Brief Introduction to Transformer

Transformer [66] is a general encoder-decoder framework built upon self-attention and encoder-decoder attention that achieves state-of-the-art results on many language generation [23, 72, 75] and understanding tasks [14, 81]. It is a stacked architecture with several blocks. Each block composes two or three basic modules:

A Self-Attention Module. Self-attention is commonly used to relate different positions of a single sequence and generate a weighted representation of its inputs. Formally, given an I -element set $\mathcal{X} = \{x_1, x_2, \dots, x_I\}$, the self-attention is defined as:

$$\text{SelfAttn}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}}\right)\mathbf{V}, \quad (1)$$

where $\mathbf{Q}, \mathbf{K}, \mathbf{V}$ (*query, key, value*) are all from \mathcal{X} by a linear projection, and d is the dimension of hidden states. Note that, the self-attention module can be employed in either an encoder or a decoder and is typically associated with a residual connection and layer normalization LN [1], resulting in the final output $\text{LN}(\mathcal{X} + \text{SelfAttn}(\mathbf{Q}, \mathbf{K}, \mathbf{V}))$.

A Feed-Forward Layer. Feed-forward layer is analogous to activated fully-connected layer, that provides non-linearity to the model. For any $i \in \{1, \dots, I\}$, $\text{FFN}(x_i) = w_2 \max(w_1 x_i + b_1, 0) + b_2$, where the w 's and b 's are the parameters to be learned. Similarly, this layer is followed with a residual connection and layer normalization.

An Optional Encoder-Decoder Attention Module. To realize cross-lingual (*i.e.*, source-to-target) learning, Vaswani *et al.* [66] introduce an encoder-decoder attention that appears in the decoder only. Formally, let $\{z_1, z_2, \dots, z_I\}$ be the output of the last layer in the encoder, and \mathbf{H} is the preceding hidden state in the decoder. The encoder-decoder attention is implemented similar to Equation (1), with the difference lies in that the *query* (\mathbf{Q}) comes from \mathbf{H} , the *key* (\mathbf{K}) and *value* (\mathbf{V}) come from $\{z_1, z_2, \dots, z_I\}$. Eventually, this

module outputs $\text{LN}(\mathbf{H} + \text{EncDecAttn}(\mathbf{Q}, \mathbf{K}, \mathbf{V}))$. More details can be found in the original paper [66].

3. Dense Interaction Learning

In this section, we will give a detailed introduction to our proposed DenseIL. Our framework has a hybrid architecture and is composed of a CNN encoder and a Dense Interaction (DI) decoder. In general, with the assistance of the proposed Dense Attention, the DenseIL can enable the decoder densely attend to intermediate CNN features, naturally forming a multi-grained and interacted representation for each video clip. In the following, we first introduce each component we used in details and then give three variants for the overall architecture as demonstrated in Figure 2.

3.1. CNN Encoder

In contrast to the vanilla Transformer that employs self-attention to construct hidden features for the source inputs, we use convolution-based transformation to generate hidden features for the inputs due to both its efficiency and high performance (*e.g.*, translation equivariance and locality).

Feature Extraction. Given a set of sampled video frames $\mathcal{X} = \{x_1, x_2, \dots, x_I\}$ with length I , the CNN encoder extracts the hidden spatial features block by block, where each block can be an arbitrary CNN structure (*e.g.*, Res-Block [21], Dense-Block [31], *etc.*). We denote the spatial features in the l -th block as $\{z_1^l, z_2^l, \dots, z_I^l\}$. In order to fully utilize the previously accumulated experience of pre-training [87], we make minimal changes on the CNN architecture. Therefore, our CNN encoder can be initialized with ImageNet-pretrained parameters [13], endowing it with more power of robust representation.

Horizontal Partition. Part-based re-ID model has enjoyed rich success in person re-ID [65, 63, 38, 18], where each input sample is partitioned into patches by predefined priori knowledge or external supervision. The spirit of the part-based schemes lay on utilizing part-level features to provide more discriminative representation to distinguish the person-of-interest from others. Inspired by this, we horizontally divide the spatial features $\{z_1, z_2, \dots, z_I\}$ produced by CNN encoder into P feature patches, and then perform average pooling on the divided feature to build a part-level feature vector $z_i^p \in \mathbb{R}^{1 \times d}$ for each patch, where $p \in \{1, \dots, P\}$, and d is the number of channels of the spatial features generated by CNN encoder. Eventually, we stack all feature vectors and obtain $\mathbf{Z} = [z_1^1; \dots; z_I^P]$, where $\mathbf{Z} \in \mathbb{R}^{IP \times d}$ and $[\cdot]$ denotes the concatenation operation.

3.2. Dense Interaction Decoder

As the CNN encoder generates the spatial features for each video clip, we further adopt the Dense Interaction (DI)

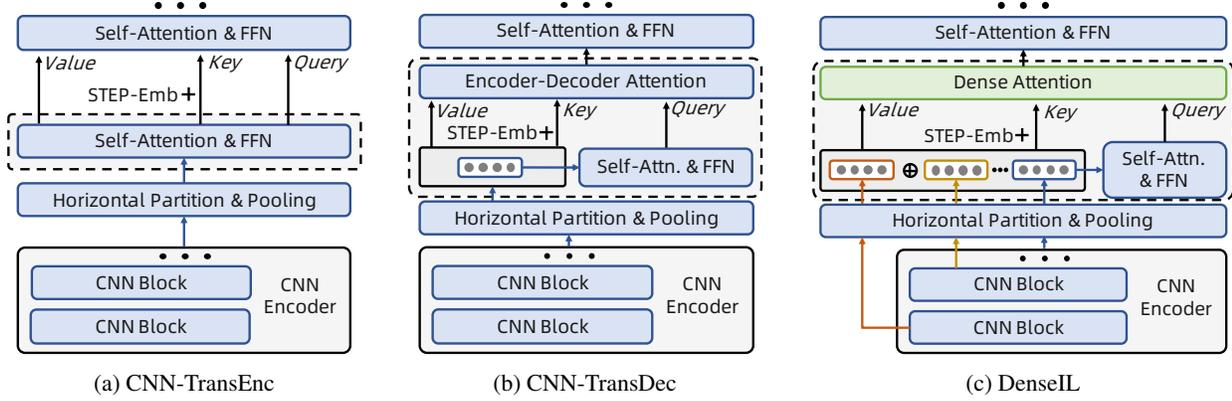


Figure 2: The proposed three model variants for the video-based person re-ID task. (a) The decoder only consists of self-attention (is equivalent to the encoder of vanilla Transformer). (b) The decoder contains both self-attention and encoder-decoder attention (is equivalent to the decoder of vanilla Transformer). (c) Our DI decoder involves self-attention and the proposed Dense Attention (The \oplus denotes the concatenation operation). All schemes are equipped with our proposed STEP-Emb. We omit the layer normalization for simplicity.

decoder to model the long-range context, especially for excavating intra-frame and inter-frame relationships. In general, our DI decoder follows the design philosophy of the self-attention module and feed-forward layer in Transformer, but differs in the encoder-decoder attention module, which is substituted for the proposed Dense Attention module. The main difference between the vanilla encoder-decoder attention and our Dense Attention is that ours allows the decoder to additionally attend to intermediate fine-grained CNN features, forming a hybrid dense connection between the multi-grained CNN and the attention features.

Self-Attention. Our self-attention layer is developed with a focus on modeling the interaction between feature patches by reusing the innovations explored in prior works [52, 66, 44]. For the stacked feature vectors Z , we conduct self-attention in Equation (1), where

$$Q, K, V = ZW_q, ZW_k, ZW_v, \quad (2)$$

$Z \in \mathbb{R}^{IP \times d}$, and W_q, W_k, W_v are three learnable matrices to project Z into different spaces. In particular, we implement $\text{SelfAttn}(\cdot)$ with the multi-head attention as that proposed in Vaswani *et al.* [66]. Our self-attention module is also associated with a residual connection and layer normalization $\text{LN}(\cdot)$ as described in Section 2.2.

Recent study [76] also shows that Pre-LN (put the $\text{LN}(\cdot)$ inside the residual connection) demonstrates stronger stability than the original Post-LN (place the $\text{LN}(\cdot)$ between the residual blocks). In this paper, we thereby adopt Pre-LN structure in our DI decoder.

Feed-Forward Layer. We follow the spirit of the feed-forward layer $\text{FFN}(\cdot)$ of vanilla Transformer to endow the decoder with non-linearity.

The Proposed Dense Attention. Residual networks [21, 31] build skip connections that can easily accumulate the features from the previous layer to the next layer, and can achieve great performance in a wide range of tasks. Inspired by this, we propose the Dense Attention, an operation that shares similar insights to the residual networks but realizes skip connections with attention mechanism, instead of summation [21] or concatenation [31].

Concretely, the Dense Attention simultaneously attends to the intermediate features $Z^l = \{z_1^l, z_2^l, \dots, z_I^l\}$ extracted by CNN, where $l \in \{1, \dots, L\}$, and the high-level features $H^r = \{h_1^r, h_2^r, \dots, h_{IP}^r\}$ generated from the preceding self-attention module (and the feed-forward layer) in the r -th block, where $r \in \{1, \dots, R\}$. For the intermediate CNN features Z^l , which retain the fine-grained spatial information of inputs, we successively perform horizontal partition and an average pooling operation, denoted as $\text{PPool}(\cdot)$, to compress each partition into a feature vector before fed into Dense Attention. On the other hand, the features H^r are directly utilized, which represents the relationship between partitioned patches across frames. More precisely, our Dense Attention $\text{DenseAttn}(Q, K, V)$ is executed in a similar way as Equation (1), with modifications that the *query* (Q) comes from H^r , the *key* (K) and *value* (V) comes from the concatenation of the pooled intermediate CNN features $\text{PPool}(Z^l)$ and the output of self-attention module H^r :

$$\begin{aligned} V^r &= [\text{PPool}(Z^1), \dots, \text{PPool}(Z^L), H^r] W_v^r, \\ K^r &= [\text{PPool}(Z^1), \dots, \text{PPool}(Z^L), H^r] W_k^r, \\ Q^r &= H^r W_q^r, \end{aligned} \quad (3)$$

where $[\cdot]$ represents the concatenation operation. The detailed distinction among self-attention, encoder-decoder at-

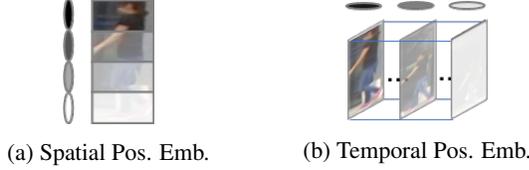


Figure 3: The spatial and temporal positional embedding contribute to the proposed STEP-Emb together.

tention and Dense Attention will be elaborated in Section 3.4.

3.3. Spatio-Temporal Positional Embedding

Since attention mechanism has no recurrent operation like RNN [25] or convolution like CNN (*i.e.*, permutation-variant), we need explicitly introduce the absolute or relative position of inputs to the model. Thus, we equip our DI decoder with positional embedding as demonstrated in Figure 3. The positional embedding is twofold: a spatial embedding and a temporal one. Both of them contribute to the final Spatial-Temporal Positional Embedding (STEP-Emb).

Spatial Positional Embedding. As illustrated in Figure 3a, the spatial positional embedding is constructed based on the partitioned spatial features, denoted as $z_i = \{z_i^1, z_i^2, \dots, z_i^P\}$ for i -th frame. We use sinusoidal embedding function to format embedding vectors [19, 66]: $\text{SpatialPos}(p, j)$ equals to $\sin(p/10000^{j/d})$ if j is even, and equals to $\cos(p/10000^{j/d})$ otherwise. Here $p \in \{1, \dots, P\}$ is the index of the partition, and $j \in \{1, \dots, d\}$ is the index of the hidden dimension. For simplicity, we denote $s^p = \text{SpatialPos}(p, j)$.

Temporal Positional Embedding. Compared with image-based recognition, video sequences additionally leak motion context for identifying pedestrians. However, Attention-based architectures (*i.e.*, self-attention) are naturally permutation-invariant [66]. Thus, they are incapable of explicitly modeling chronological relation between video frames, leading to a disruption of the intrinsic temporal structure. To tackle this problem, apart from the spatial positional embedding, we also build a temporal positional embedding to imply chronological order of inputs, as shown Figure 3b. The temporal positional embedding can be formulated as: $\text{TemporalPos}(i, j)$ equals to $\sin(i/10000^{j/d})$ if j is even, and equals to $\cos(i/10000^{j/d})$ otherwise. Here $i \in \{1, \dots, I\}$ is the index of input frame. For simplicity, we denote $t_i = \text{TemporalPos}(i, j)$.

Spatial-Temporal Positional Embedding. Based on the above two strategies, we derive the final Spatio-TEmporal Positional Embedding (STEP-Emb), which is constructed by the summation of the spatial and the temporal embedding. More precisely, the STEP-Emb for each feature patch z_i^p can be represented as: $e_i^p = s^p + t_i$. Intuitively, the STEP-Emb

improves the power of capturing long range interaction by enhancing the structural information between partitioned patches. We will demonstrate its ability in Section 4.4.

3.4. Overall Architecture

To dive deeply into the CNN-Attention hybrid structure, we introduce three model variants for the overall architecture. Figure 2 demonstrates the detailed configurations. For all the three variants, the components in the dashed boxes can be regarded as the basic building blocks to stack up. The final outputs of the decoder are processed by a batch normalization layer and a non-bias classifier layer.

CNN encoder & Vanilla Transformer Encoder (CNN-TransEnc). We directly cascade the CNN encoder, and the vanilla Transformer encoder, which mainly comprises self-attention module ($\text{SelfAttn}(\cdot)$), to model interactions between feature partitions across frames. Therefore, the *query*, *key* and *value* are all from the preceding self-attention module. The resulting architecture is illustrated in Figure 2a.

CNN encoder & Vanilla Transformer Decoder (CNN-TransDec). We employ vanilla Transformer decoder combined with our CNN encoder as another model variant. As shown in Figure 2b, compared with CNN-TransEnc, CNN-TransDec additionally includes the encoder-decoder attention ($\text{EncDecAttn}(\cdot)$) to pay closer attention to the highest-level spatial features generated by the CNN encoder.

CNN encoder & Our DI Decoder with Dense Attention (DenseIL). The aforementioned CNN-TransEnc and CNN-TransDec only take the last-layer (the highest level) representations from the CNN encoder as inputs. As a comparison, our DenseIL (see Figure 2c) is equipped with the proposed Dense Attention, who is able to simultaneously attend to multi-scale intermediate spatial features in stacked CNN blocks and the corresponding interaction modeled by self-attention module.

Loss Function. We adopt the simplest cross-entropy loss and batch triplet loss [24] for a fair comparison with the previous works [61, 26, 20]. The two loss functions are calculated on the global average-pooled sequence-level feature vectors generated by the DI decoder.

4. Experimental Results

4.1. Datasets and Evaluation Protocol

Datasets. We evaluate our DenseIL on several commonly adopted video-based re-ID benchmarks, including MARS [87], DukeMTMC-VideoReID (DukeV) [57, 74] and iLIDS-VID [69].

(a) We compare Dense Attention with its counterparts.

Methods	mAP	R-1	R-5
Baseline	82.1	87.3	95.6
CNN-TransEnc	85.7	89.4	96.6
CNN-TransDec	85.8	90.2	96.5
w/ Dense Concat.	86.4	89.8	96.8
w/ Dense Sum.	86.2	89.9	96.7
w/ Dense Attn.	87.0	90.8	97.1

(d) Dense Attention with multi-grained CNN features. (H^r is omitted for simplicity).

key / value	mAP	R-1	R-5
Baseline	82.1	87.3	95.6
Z^4 (CNN-TransDec)	85.8	90.2	96.5
$Z^4 + Z^3$	86.7	90.6	97.1
$Z^4 + Z^3 + Z^2$	87.0	90.8	97.1
$Z^4 + Z^3 + Z^2 + Z^1$	86.9	90.5	97.1

(b) DenseIL with / without positional embedding.

Methods	mAP	R-1	R-5
Baseline	82.1	87.3	95.6
DenseIL w/o Pos. Emb.	86.3	89.9	97.0
w/ Spatial-Emb.	86.6	90.5	97.1
w/ Temporal-Emb.	86.6	90.2	97.1
w/ STEP-Emb.	87.0	90.8	97.1

(e) Different number of dimensions of hidden states with fixed blocks $R = 4$.

# Dims	GFs	mAP	R-1	R-5
Baseline	16.39	82.1	87.3	95.6
$d = 256$	16.52	86.6	90.1	97.1
$d = 512$	16.75	86.5	90.3	97.2
$d = 1024$	17.48	86.8	90.6	97.0
$d = 2048$	20.06	87.0	90.8	97.1

(c) Different number of blocks (R) with fixed $d = 2048$.

# Blocks	GFs	mAP	R-1	R-5
Baseline ($R = 0$)	16.39	82.1	87.3	95.6
$R = 1$	17.43	84.2	88.2	96.3
$R = 2$	18.31	86.3	89.6	96.9
$R = 3$	19.18	86.6	90.1	97.1
$R = 4$	20.06	87.0	90.8	97.1

(f) We vary the number of partitions for each frame.

# Partitions	mAP	R-1	R-5
Baseline	82.1	87.3	95.6
$P = 1$	86.1	89.7	96.9
$P = 2$	86.8	90.5	97.1
$P = 4$	87.0	90.8	97.1
$P = 8$	86.6	90.5	97.0

Table 1: Ablation study on MARS dataset. GFs means GFLOPs. More details are explained in the text.

Evaluation Protocol. Following the common practices [88, 87, 41, 80], we resort to both Cumulative Matching Characteristic (CMC) curves at Rank-1 (R-1) to Rank-20 (R-20), and mean Average Precision (mAP) as evaluation metrics. We do not use re-ranking for all settings.

4.2. Implementation Details

We employ ImageNet-pretrained standard ResNet-50 as the initialization of our CNN encoder. To be comparable with the previous works [63, 45, 85, 80], we also remove the last spatial down-sampling operation in the `conv5_x` block for both the baseline and our schemes. Each input frame is resized to 256×128 with frame-level random horizontal flips [45] and sequence-level random erasing [90, 85, 9] for data augmentation. We adopt a restricted random sampling strategy [41, 45, 80] to randomly sample frames from equally divided 8 chunks for each video clip. For each training batch, we randomly sample 16 identities, each with 4 tracklets [24, 80]. We train our network for 800 epoch with Adam optimizer for both the cross-entropy loss and the triplet loss [24]. The initial learning rate is set to 0.0002 and is decayed by 10 every 100 epochs. The algorithm is implemented with PyTorch¹ and trained on a 4-GPU machine.

4.3. Study on Design Choices

Spatial-Temporal Interaction Matters. We present the performance (%) of CNN-TransEnc, CNN-TransDec and DenseIL (w/ Dense Attn.) on MARS dataset in Table 1a, from which we can observe that solely allocating self-

attention module to the decoder (CNN-TransEnc) already brings effective performance gain (+3.6% mAP), while the one additionally inserted with encoder-decoder attention (CNN-TransDec) achieves comparable results (+3.7% mAP). Such observation demonstrates the strong interaction-modeling capability of self-attention and exactly verifies one of our starting point: spatial-temporal interaction matters for video-based re-ID.

Dense Attention vs. Dense Summation/Concatenation.

Unlike ResNet [21] and DenseNet [31] that combine features through summation or concatenation, we densely attend to preceding CNN and self-attention features by our proposed Dense Attention. Accordingly, we conduct experiments on densely summing (w/ Dense Sum.) or concatenating (w/ Dense Concat.) preceding CNN features to the decoder. To be more specific, we perform partition and pooling operation on the spatial features of each CNN block, and sum / concatenate them into the output of each self-attention module. We follow the same training strategy as described in Section 4.2 for both schemes and report mAP and CMC scores on MARS dataset in Table 1a. As a result, the Dense Attention (w/ Dense Attn.) reaches the highest performance over the two counterparts (+0.6% mAP). On the one hand, compared with summation or concatenation, our Dense Attention is able to adaptively learns the preference on low-level CNN features or high-level spatial-temporal interaction. On the other hand, the features learned by the CNN and the attention module may not lie in the same manifold, and concatenating them together can result in distribution gap. Instead, our Dense Attention enable better coordination between CNN and Attention by a relatively soft connection scheme.

¹Based on: https://github.com/yuange250/not_so_strong_baseline_for_video_based_person_reID.

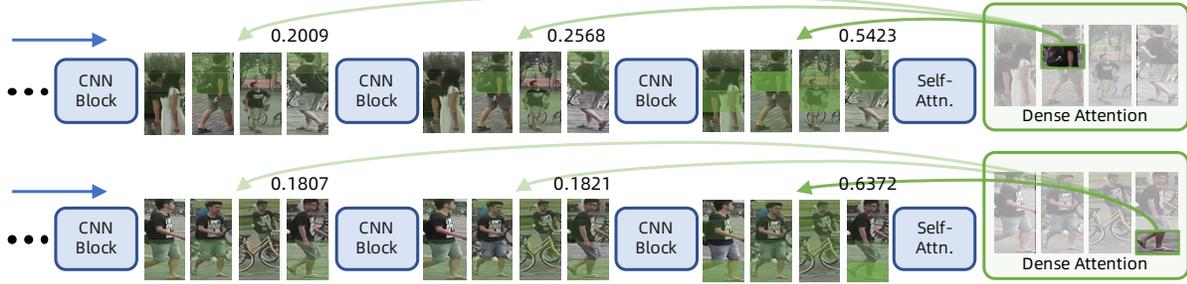


Figure 4: Visualization of the attention weights learned in Dense Attention. Darker color means higher attention weight.

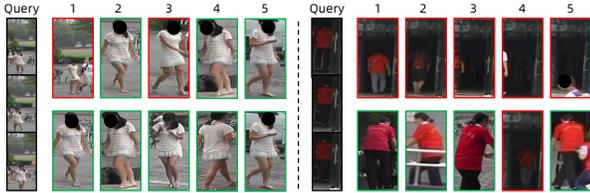


Figure 5: Visualization of re-identification results of the baseline (top) and DenseIL (bottom). The left column of each case is sampled frames of query sequence and the right five columns are top-5 retrieved sequences in the gallery set, where the item annotated with green box is correctly re-identified, and the red box denotes the wrong results.

How Dense Attention Works? In order to investigate the functional principle of our Dense Attention, we adjust the composition of its *key* and *memory* in Equation (3) and conduct corresponding experiments. More precisely, let Z^1 , Z^2 , Z^3 and Z^4 be the intermediate features generated by four CNN blocks (ResNet-50), and H^r be the features generated from the preceding self-attention module. We allow Dense Attention to attend to H^r and part of $\{Z^1, Z^2, Z^3, Z^4\}$. The results on MARS dataset are demonstrated in Table 1d. We can observe that when attending to H^r and any of $\{Z^2, Z^3, Z^4\}$, DenseIL outperforms the CNN-TransDec counterpart by a large margin. It is also interesting to see that when further attending to the lowest-level CNN feature Z^1 , the performance slightly drops (-0.1% mAP) compared with $\{Z^2, Z^3, Z^4\}$. This might be due to the fact that the lowest-level CNN feature lacks discriminative information and is not conducive to the coordination with the attention scheme in terms of video re-ID task. For a better understanding of Dense Attention, we randomly feed image samples into DenseIL and map the attention weights of Dense Attention to colors. The results are visualized in Figure 4, which suggests that the Dense Attention indeed cares about multi-grained CNN features simultaneously, while gives relatively more concern on the highest-level one. We also demonstrate the retrieval results in Figure 5. We can observe that, in the left case, although there exists misalignment, movement and occlusion in the query respectively,

our scheme is still able to match the person-of-interest accurately. In contrast, the baseline model misses the sequences of the same identity. In the right case, the baseline model re-identifies the query incorrectly due to ignoring the fine-grained information between visually similar identities. In contrast, DenseIL captures the contour and the fine-grained characters on her back, yielding a satisfactory re-ID result.

STEP-Emb Tackles Permutation-Invariant Issue. The philosophy behind STEP-Emb is to empower the DI decoder to capture long-range interaction by enhancing structural information between partitioned patches with embeddings. To verify its functionality, we remove the STEP-Emb (DenseIL w/o Pos. Emb.) and demonstrate its performance (%) on MARS dataset in Table 1b. It can be easily observed that employing either the Spatial Positional Embedding (w/ Spatial-Emb.) or the Temporal Positional Embedding (w/ Temporal-Emb) undeniably boosts the person re-ID performance (+0.3% mAP). As the joint one, STEP-Emb, achieves the best performance consistently on all metrics. This suggests that STEP-Emb is indispensable for spatial-temporal positional modeling in our DI decoder.

4.4. Study on Model Variations

Deeper DI Decoder Improves the Performance. As described in Section 3, our DI decoder is stacked with R basic building blocks. Correspondingly, we vary different number of blocks to investigate how our model performs. Table 1c illustrates the performance (%) on MARS dataset with $R = 0/1/2/3/4$ blocks. We can observe a consistent performance improvement as we increase the number of blocks in DI decoder.

Wider DI Decoder Improves the Performance. We vary dimension of hidden states d to discover how the width of DI decoder affects the re-ID performance. Table 1e demonstrates the corresponding results, which tell that wider DI decoder indeed shows larger capability, resulting consistent performance gain on MARS dataset.

Horizontal Partition Helps Our DI Decoder to Learn Discriminative Features. Part-based techniques for person

Methods	Proc.	Backbone	MARS				DukeV				iLIDS-VID	
			mAP	R-1	R-5	R-20	mAP	R-1	R-5	R-10	R-1	R-5
STAN [41]	CVPR18	Res50	65.8	82.3	-	-	-	-	-	-	80.2	-
Snippet [8]	CVPR18	Res50	76.1	86.3	94.7	98.2	-	-	-	-	85.4	96.7
STA [17]	AAAI19	Res50	80.8	86.3	95.7	-	94.9	96.2	99.3	99.6	-	-
ADFD [86]	CVPR19	Res50	78.2	87.0	95.4	98.7	-	-	-	-	86.3	97.4
VRSTC [27]	CVPR19	Res50	82.3	88.5	96.5	-	93.5	95.0	99.1	99.4	83.4	95.5
GLTR [39]	ICCV19	Res50	78.5	87.0	95.8	98.2	93.7	96.3	99.3	-	86.0	98.0
COSAM [61]	ICCV19	SE-Res50	79.9	84.9	95.5	97.9	94.1	95.4	99.3	-	79.6	95.3
STE-NVAN [45]	BMVC19	Res50-NL	81.2	88.9	-	-	93.5	95.2	-	-	-	-
MG-RAFA [85]	CVPR20	Res50	85.9	88.8	97.0	98.5	-	-	-	-	88.6	98.0
MGH [79]	CVPR20	Res50-NL	85.8	90.0	96.7	98.5	-	-	-	-	85.6	97.1
STGCN [80]	CVPR20	Res50	83.7	90.0	96.4	98.3	95.7	97.3	99.3	-	-	-
TCLNet [26]	ECCV20	Res50-TCL	85.1	89.8	-	-	96.2	96.9	-	-	86.6	-
AP3D [20]	ECCV20	AP3D	85.1	90.1	-	-	95.6	96.3	-	-	86.7	-
AFA [9]	ECCV20	Res50	82.9	90.2	96.6	-	95.4	97.2	99.4	99.7	88.5	96.8
Ours	-	Res50	87.0	90.8	97.1	98.8	97.1	97.6	99.7	99.9	92.0	98.0

Table 2: We compare the DenseIL with state-of-the-art results.

re-ID has remained attractive in the last few years [63, 17, 79, 80, 83], we here investigate how horizontal partition influence the learning of DenseIL. The results are illustrated in Table 1f, from which we can conclude that the best performance is reached when we partition each spatial feature into four patches. Interestingly but not surprisingly, training with larger number of partitions does not achieve higher result (-0.4% mAP). This might be due to more partitions amplify the misalignment issue existing in video-based person re-ID.

Our DI Decoder is not Limited to Specific CNN Backbone. For example, DenseIL is able to boost the mAP of DenseNet-121 baseline from 82.5% to 86.7% on MARS dataset, demonstrating great generalization ability on different CNN basic structure.

4.5. Computational Complexity

For all settings, we compute the theoretical GFLOPS with the tool `compute_flops.py`² from DETR [5]. From Table 1c and 1e we can observe that, the deeper and the wider the DI decoder is, the higher computational complexity we have. However, such issue can be alleviated with the recently developed high efficiency Transformer variants [35, 94]. In addition, for the lightweight version in Table 1e, *e.g.*, $d = 256$, which brings negligible computational overhead, still outperforms the baseline by 4.5% mAP.

4.6. Comparison with State-of-the-Art Results

We compare our proposed DenseIL with the best results reported in recent literatures in Table 2. We can see that our

²<https://gist.github.com/fmassa/c0fbb9fe7bf53b533b5cc241f5c8234c>

method achieves significant improvement over all the competitive state-of-the-arts, including the optical flow-based [9], graph-based [80, 79], 3D CNN-based [40, 20] and attention-based [17, 27, 39, 61, 26] approaches, in terms of both the mAP and CMC metrics. For example, our method surpasses the latest optical flow-based scheme [9] by 4.1% mAP, outperforms the latest graph-based scheme [80] by 3.3% mAP on MARS dataset. For attention-based approaches [39, 26], our method is also superior to all of them by a large margin. We argue that the performance gain comes from the usage of both the multi-grained cues and the spatial-temporal interaction, which are not fully exploited in previous works.

5. Conclusion

In this paper, we successfully leverage Dense Interaction Learning (DenseIL) to alleviate the difficulties of multi-grained spatial-temporal interaction modeling for video-based person re-ID. Specifically, by incorporating the proposed Dense Attention and STEP-Emb, we let the DI decoder densely attends to intermediate CNN features and generates an intrinsically multi-grained representation for each video clip. Experimental results demonstrate that our approach surpasses all previous methods.

For future works, it is interesting to apply Dense Interaction Learning to more video understanding tasks, such as action recognition, video captioning, *etc.*

Acknowledgements. This work was supported in part by the National Key Research and Development Program of China 2018AAA0101400 and NSFC Grant U1908209, 61632001, 62021001. We thank all the anonymous reviewers for their valuable comments on our paper.

References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. [2](#), [3](#)
- [2] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015*, 2015. [2](#), [3](#)
- [3] Irwan Bello, Barret Zoph, Ashish Vaswani, Jonathon Shlens, and Quoc V Le. Attention augmented convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3286–3295, 2019. [2](#)
- [4] Yue Cao, Jiarui Xu, Stephen Lin, Fangyun Wei, and Han Hu. Gcnet: Non-local networks meet squeeze-excitation networks and beyond. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019. [2](#)
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision*, 2020. [2](#), [8](#)
- [6] William Chan, Navdeep Jaitly, Quoc Le, and Oriol Vinyals. Listen, attend and spell: A neural network for large vocabulary conversational speech recognition. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4960–4964. IEEE, 2016. [3](#)
- [7] Xiaobin Chang, Timothy M Hospedales, and Tao Xiang. Multi-level factorisation net for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2109–2118, 2018. [1](#)
- [8] Dapeng Chen, Hongsheng Li, Tong Xiao, Shuai Yi, and Xiang Wang. Video person re-identification with competitive snippet-similarity aggregation and co-attentive snippet embedding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1169–1178, 2018. [8](#)
- [9] Guangyi Chen, Yongming Rao, Jiwen Lu, and Jie Zhou. Temporal coherence or temporal motion: Which is more critical for video-based person re-identification? In *European Conference on Computer Vision*, 2020. [1](#), [2](#), [6](#), [8](#)
- [10] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, 2014. [3](#)
- [11] Dahjung Chung, Khalid Tahboub, and Edward J Delp. A two stream siamese convolutional neural network for person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1983–1991, 2017. [2](#)
- [12] Marcella Cornia, Matteo Stefanini, Lorenzo Baraldi, and Rita Cucchiara. Meshed-memory transformer for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10578–10587, 2020. [2](#)
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [3](#)
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 2019. [3](#)
- [15] Cheng-Yang Fu, Wei Liu, Ananth Ranga, Amrith Tyagi, and Alexander C Berg. Dssd: Deconvolutional single shot detector. *arXiv preprint arXiv:1701.06659*, 2017. [3](#)
- [16] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3146–3154, 2019. [2](#)
- [17] Yang Fu, Xiaoyang Wang, Yunchao Wei, and Thomas Huang. Sta: Spatial-temporal attention for large-scale video-based person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8287–8294, 2019. [2](#), [8](#)
- [18] Yang Fu, Yunchao Wei, Yuqian Zhou, Honghui Shi, Gao Huang, Xinchao Wang, Zhiqiang Yao, and Thomas Huang. Horizontal pyramid matching for person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8295–8302, 2019. [3](#)
- [19] Jiatao Gu, James Bradbury, Caiming Xiong, Victor OK Li, and Richard Socher. Non-autoregressive neural machine translation. In *International Conference on Learning Representations*, 2018. [5](#)
- [20] Xinqian Gu, Hong Chang, Bingpeng Ma, Hongkai Zhang, and Xilin Chen. Appearance-preserving 3d convolution for video-based person re-identification. In *European Conference on Computer Vision*. Springer, 2020. [2](#), [5](#), [8](#)
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [2](#), [3](#), [4](#), [6](#)
- [22] Tianyu He, Xu Shen, Jianqiang Huang, Zhibo Chen, and Xian-Sheng Hua. Partial person re-identification with part-part correspondence learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9105–9115, 2021. [1](#)
- [23] Tianyu He, Xu Tan, Yingce Xia, Di He, Tao Qin, Zhibo Chen, and Tie-Yan Liu. Layer-wise coordination between encoder and decoder for neural machine translation. In *Advances in Neural Information Processing Systems*, pages 7944–7954, 2018. [3](#)
- [24] Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*, 2017. [5](#), [6](#)
- [25] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. [5](#)
- [26] Ruibing Hou, Hong Chang, Bingpeng Ma, Shiguang Shan, and Xilin Chen. Temporal complementary learning for video person re-identification. 2020. [2](#), [5](#), [8](#)
- [27] Ruibing Hou, Bingpeng Ma, Hong Chang, Xinqian Gu, Shiguang Shan, and Xilin Chen. Vrstc: Occlusion-free video

- person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7183–7192, 2019. 2, 8
- [28] Han Hu, Zheng Zhang, Zhenda Xie, and Stephen Lin. Local relation networks for image recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3464–3473, 2019. 2
- [29] Jie Hu, Li Shen, Samuel Albanie, Gang Sun, and Andrea Vedaldi. Gather-excite: Exploiting feature context in convolutional neural networks. In *Advances in neural information processing systems*, pages 9401–9411, 2018. 2
- [30] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018. 2
- [31] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. 2, 3, 4, 6
- [32] Xin Jin, Cuiling Lan, Wenjun Zeng, Zhibo Chen, and Li Zhang. Style normalization and restitution for generalizable person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3143–3152, 2020. 1
- [33] Xin Jin, Cuiling Lan, Wenjun Zeng, Guoqiang Wei, and Zhibo Chen. Semantics-aligned representation learning for person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11173–11180, 2020. 1
- [34] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2016. 1
- [35] Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. In *International Conference on Learning Representations*, 2019. 8
- [36] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. 1
- [37] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 2
- [38] Dangwei Li, Xiaotang Chen, Zhang Zhang, and Kaiqi Huang. Learning deep context-aware features over body and latent parts for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 384–393, 2017. 3
- [39] Jianing Li, Jingdong Wang, Qi Tian, Wen Gao, and Shiliang Zhang. Global-local temporal representations for video person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3958–3967, 2019. 2, 8
- [40] Jianing Li, Shiliang Zhang, and Tiejun Huang. Multi-scale 3d convolution network for video based person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8618–8625, 2019. 1, 2, 8
- [41] Shuang Li, Slawomir Bak, Peter Carr, and Xiaogang Wang. Diversity regularized spatiotemporal attention for video-based person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 369–378, 2018. 1, 2, 6, 8
- [42] Wei Li, Rui Zhao, Tong Xiao, and Xiaogang Wang. Deepreid: Deep filter pairing neural network for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 152–159, 2014. 1
- [43] Xiang Li, Wei-Shi Zheng, Xiaojuan Wang, Tao Xiang, and Shaogang Gong. Multi-scale learning for low-resolution person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3765–3773, 2015. 1
- [44] Zhouhan Lin, Minwei Feng, Cicero Nogueira dos Santos, Mo Yu, Bing Xiang, Bowen Zhou, and Yoshua Bengio. A structured self-attentive sentence embedding. In *International Conference on Learning Representations*, 2017. 4
- [45] Chih-Ting Liu, Chih-Wei Wu, Yu-Chiang Frank Wang, and Shao-Yi Chien. Spatially and temporally efficient non-local attention network for video-based person re-identification. In *BMVC*, 2019. 2, 6, 8
- [46] Yu Liu, Junjie Yan, and Wanli Ouyang. Quality aware network for set to set recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5790–5799, 2017. 2
- [47] Yiheng Liu, Zhenxun Yuan, Wengang Zhou, and Houqiang Li. Spatial and temporal mutual promotion for video-based person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8786–8793, 2019. 2
- [48] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 3
- [49] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vibert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems*, pages 13–23, 2019. 2
- [50] Xiaojiao Mao, Chunhua Shen, and Yu-Bin Yang. Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections. In *Advances in neural information processing systems*, pages 2802–2810, 2016. 3
- [51] Niall McLaughlin, Jesus Martinez Del Rincon, and Paul Miller. Recurrent convolutional network for video-based person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1325–1334, 2016. 2
- [52] Ankur Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. A decomposable attention model for natural language inference. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2249–2255, 2016. 4
- [53] Niki Parmar, Prajit Ramachandran, Ashish Vaswani, Irwan Bello, Anselm Levskaya, and Jon Shlens. Stand-alone self-attention in vision models. In *Advances in Neural Information Processing Systems*, pages 68–80, 2019. 2
- [54] Xuelin Qian, Yanwei Fu, Yu-Gang Jiang, Tao Xiang, and Xiangyang Xue. Multi-scale deep learning architectures for

- person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5399–5408, 2017. 1
- [55] Zhaofan Qiu, Ting Yao, and Tao Mei. Learning spatio-temporal representation with pseudo-3d residual networks. In *proceedings of the IEEE International Conference on Computer Vision*, pages 5533–5541, 2017. 2
- [56] Wasifur Rahman, Md Kamrul Hasan, Sangwu Lee, AmirAli Bagher Zadeh, Chengfeng Mao, Louis-Philippe Morency, and Ehsan Hoque. Integrating multimodal information in large pretrained transformers. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2359–2369, 2020. 2
- [57] Ergys Ristani, Francesco Solera, Roger Zou, Rita Cucchiara, and Carlo Tomasi. Performance measures and a data set for multi-target, multi-camera tracking. In *European Conference on Computer Vision*, pages 17–35. Springer, 2016. 5
- [58] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 3
- [59] Chi Su, Jianing Li, Shiliang Zhang, Junliang Xing, Wen Gao, and Qi Tian. Pose-driven deep convolutional model for person re-identification. In *Proceedings of the IEEE international conference on computer vision*, pages 3960–3969, 2017. 1
- [60] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vi-bert: Pre-training of generic visual-linguistic representations. In *International Conference on Learning Representations*, 2019. 2
- [61] Arulkumar Subramaniam, Athira Nambiar, and Anurag Mittal. Co-segmentation inspired attention networks for video-based person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 562–572, 2019. 2, 5, 8
- [62] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7464–7473, 2019. 2
- [63] Yifan Sun, Liang Zheng, Yi Yang, Qi Tian, and Shengjin Wang. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 480–496, 2018. 1, 3, 6, 8
- [64] Tong Tong, Gen Li, Xiejie Liu, and Qinquan Gao. Image super-resolution using dense skip connections. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4799–4807, 2017. 3
- [65] Rahul Rama Variar, Bing Shuai, Jiwen Lu, Dong Xu, and Gang Wang. A siamese long short-term memory architecture for human re-identification. In *European conference on computer vision*, pages 135–153. Springer, 2016. 3
- [66] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. 2, 3, 4, 5
- [67] Guanshuo Wang, Yufeng Yuan, Xiong Chen, Jiwei Li, and Xi Zhou. Learning discriminative features with multiple granularities for person re-identification. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 274–282, 2018. 1
- [68] Huiyu Wang, Yukun Zhu, Bradley Green, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Axial-deeplab: Stand-alone axial-attention for panoptic segmentation. In *Proceedings of the European Conference on Computer Vision*, 2020. 2
- [69] Taiqing Wang, Shaogang Gong, Xiatian Zhu, and Shengjin Wang. Person re-identification by video ranking. In *European conference on computer vision*, pages 688–703. Springer, 2014. 5
- [70] Xiaogang Wang. Intelligent multi-camera video surveillance: A review. *Pattern recognition letters*, 34(1):3–19, 2013. 1
- [71] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018. 2
- [72] Yiren Wang, Yingce Xia, Tianyu He, Fei Tian, Tao Qin, ChengXiang Zhai, and Tie-Yan Liu. Multi-agent dual learning. In *Proceedings of the International Conference on Learning Representations (ICLR) 2019*, 2019. 3
- [73] Longhui Wei, Shiliang Zhang, Hantao Yao, Wen Gao, and Qi Tian. Glad: Global-local-alignment descriptor for pedestrian retrieval. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 420–428, 2017. 1
- [74] Yu Wu, Yutian Lin, Xuanyi Dong, Yan Yan, Wanli Ouyang, and Yi Yang. Exploit the unknown gradually: One-shot video-based person re-identification by stepwise learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5177–5186, 2018. 2, 5
- [75] Yingce Xia, Tianyu He, Xu Tan, Fei Tian, Di He, and Tao Qin. Tied transformers: Neural machine translation with shared encoder and decoder. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5466–5473, 2019. 3
- [76] Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tie-Yan Liu. On layer normalization in the transformer architecture. In *Proceedings of the International Conference on Machine Learning*, 2020. 4
- [77] Shuangjie Xu, Yu Cheng, Kang Gu, Yang Yang, Shiyu Chang, and Pan Zhou. Jointly attentive spatial-temporal pooling networks for video-based person re-identification. In *Proceedings of the IEEE international conference on computer vision*, pages 4733–4742, 2017. 2
- [78] Yichao Yan, Bingbing Ni, Zhichao Song, Chao Ma, Yan Yan, and Xiaokang Yang. Person re-identification via recurrent feature aggregation. In *European Conference on Computer Vision*, pages 701–716. Springer, 2016. 2
- [79] Yichao Yan, Jie Qin, Jiabin Chen, Li Liu, Fan Zhu, Ying Tai, and Ling Shao. Learning multi-granular hypergraphs for video-based person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2899–2908, 2020. 1, 2, 8

- [80] Jinrui Yang, Wei-Shi Zheng, Qize Yang, Ying-Cong Chen, and Qi Tian. Spatial-temporal graph convolutional network for video-based person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3289–3299, 2020. 1, 2, 6, 8
- [81] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*, pages 5753–5763, 2019. 3
- [82] Yanhong Zeng, Jianlong Fu, and Hongyang Chao. Learning joint spatial-temporal transformations for video inpainting. In *European Conference on Computer Vision*, pages 528–543. Springer, 2020. 2
- [83] Dong Zhang, Hanwang Zhang, Jinhui Tang, Meng Wang, Xiansheng Hua, and Qianru Sun. Feature pyramid transformer. In *Proceedings of the European Conference on Computer Vision*, 2020. 2, 8
- [84] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *International Conference on Machine Learning*, pages 7354–7363. PMLR, 2019. 2
- [85] Zhizheng Zhang, Cuiling Lan, Wenjun Zeng, and Zhibo Chen. Multi-granularity reference-aided attentive feature aggregation for video-based person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10407–10416, 2020. 2, 6, 8
- [86] Yiru Zhao, Xu Shen, Zhongming Jin, Hongtao Lu, and Xiansheng Hua. Attribute-driven feature disentangling and temporal aggregation for video person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4913–4922, 2019. 8
- [87] Liang Zheng, Zhi Bie, Yifan Sun, Jingdong Wang, Chi Su, Shengjin Wang, and Qi Tian. Mars: A video benchmark for large-scale person re-identification. In *European Conference on Computer Vision*, pages 868–884. Springer, 2016. 1, 2, 3, 5, 6
- [88] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *Proceedings of the IEEE international conference on computer vision*, pages 1116–1124, 2015. 6
- [89] Liang Zheng, Yi Yang, and Qi Tian. Sift meets cnn: A decade survey of instance retrieval. *IEEE transactions on pattern analysis and machine intelligence*, 40(5):1224–1244, 2017. 1
- [90] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random erasing data augmentation. In *AAAI*, pages 13001–13008, 2020. 6
- [91] Kaiyang Zhou, Yongxin Yang, Andrea Cavallaro, and Tao Xiang. Omni-scale feature learning for person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3702–3712, 2019. 1
- [92] Zhen Zhou, Yan Huang, Wei Wang, Liang Wang, and Tieniu Tan. See the forest for the trees: Joint spatial and temporal recurrent neural networks for video-based person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4747–4756, 2017. 1, 2
- [93] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. 3
- [94] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. 2, 8