

# TF-CLIP: Learning Text-Free CLIP for Video-Based Person Re-identification

Chenyang Yu<sup>1</sup>, Xuehu Liu<sup>2</sup>, Yingquan Wang<sup>1</sup>, Pingping Zhang<sup>3\*</sup>, Huchuan Lu<sup>1,3,4</sup>

<sup>1</sup> School of Information and Communication Engineering, Dalian University of Technology, Dalian, China

<sup>2</sup> School of Computer Science and Artificial Intelligence, Wuhan University of Technology, Wuhan, China

<sup>3</sup> School of Future Technology, School of Artificial Intelligence, Dalian University of Technology, Dalian, China

<sup>4</sup> Ningbo Institute, Dalian University of Technology, Ningbo, China

yuchenyang@mail.dlut.edu.cn; liuxuehu@whut.edu.cn; yingquan\_w95@mail.dlut.edu.cn; {zhpp, lhchuan}@dlut.edu.cn

## Abstract

Large-scale language-image pre-trained models (e.g., CLIP) have shown superior performances on many cross-modal retrieval tasks. However, the problem of transferring the knowledge learned from such models to video-based person re-identification (ReID) has barely been explored. In addition, there is a lack of decent text descriptions in current ReID benchmarks. To address these issues, in this work, we propose a novel one-stage text-free CLIP-based learning framework named TF-CLIP for video-based person ReID. More specifically, we extract the identity-specific sequence feature as the CLIP-Memory to replace the text feature. Meanwhile, we design a Sequence-Specific Prompt (SSP) module to update the CLIP-Memory online. To capture temporal information, we further propose a Temporal Memory Diffusion (TMD) module, which consists of two key components: Temporal Memory Construction (TMC) and Memory Diffusion (MD). Technically, TMC allows the frame-level memories in a sequence to communicate with each other, and to extract temporal information based on the relations within the sequence. MD further diffuses the temporal memories to each token in the original features to obtain more robust sequence features. Extensive experiments demonstrate that our proposed method shows much better results than other state-of-the-art methods on MARS, LS-VID and iLIDS-VID.

## Introduction

Video-based person Re-Identification (ReID) (Liu et al. 2021c; Xu et al. 2017; Liu et al. 2015) aims at re-identifying specific persons from videos across non-overlapping cameras. Over the past decade, various approaches have been proposed to solve this challenging task, including CNN-based methods (Zhang et al. 2020; Fu et al. 2019; Subramaniam, Nambiar, and Mittal 2019; Chen et al. 2018; Dai et al. 2019) and Transformer-based methods (Liu et al. 2021b; Zhang et al. 2021a; Wu et al. 2022; Zhu et al. 2023). However, existing works employ a unimodal framework, where the video encoders are trained to predict a fixed set of predefined labels. Considering that the ReID task is in an open-set setting, such closed-set training setting limits the robustness of these methods.

\*Corresponding author.

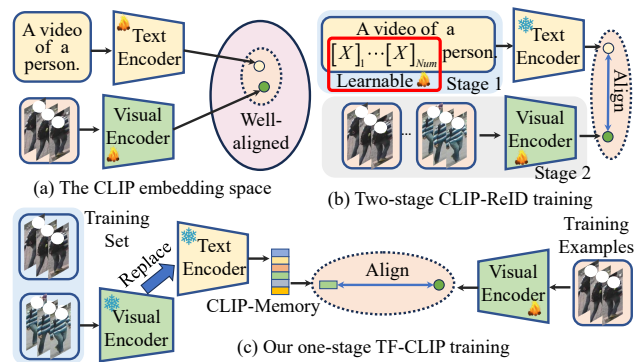


Figure 1: Comparison of CLIP-ReID and our one-stage text-free learning framework. (a) An illustration of the CLIP joint embedding space. (b) Two-stage CLIP-ReID training. (c) Our one-stage TF-CLIP training.

Recently, the outstanding Contrastive Language-Image Pre-training (CLIP) (Radford et al. 2021) has shown a great capability of learning robust representations. Different from traditional unimodal frameworks, CLIP is trained on large-scale language-image pairs in a contrastive way. Apparently, CLIP-based methods have achieved great success in video domains (Ju et al. 2022; Xu et al. 2021; Wang, Xing, and Liu 2021; Ni et al. 2022). However, the transfer and adaptation to video-based person ReID is not well explored. A major challenge is that compared with other video tasks, video sequences in video-based person ReID only have integer index labels, and there is no text labels. A straightforward solution is to annotate the natural language descriptions for existing video-based person ReID datasets. Unfortunately, it is both time consuming and labour cost to annotate suitable text descriptions for fine-grained person videos.

To overcome these issues, inspired by recent prompt engineering researches, CLIP-ReID (Li, Sun, and Li 2022) is proposed for image-based person ReID. As shown in Fig. 1 (b), CLIP-ReID fixes the text encoder and visual encoder in the first training stage, and optimizes a set of learnable text tokens to generate the label-specific text features. Then, it uses the learned text features to update the visual encoder in the second training stage. Although CLIP-ReID has achieved excellent performances, it does not consider the

temporal information in video sequences. For video-based person ReID, extracting robust temporal features is a critical step (Gao and Nevatia 2018). Thus, directly applying CLIP-ReID in video-based person ReID is sub-optimal. On the other hand, the two-stage training strategy is not elegant enough and needs hyper-parameter tuning.

To address the aforementioned issues, we propose a one-stage text-free CLIP-based person ReID framework named TF-CLIP. Specifically, as shown in Fig. 1 (c), we first use a pre-trained CLIP visual encoder for all video sequences of each identity to extract sequence features, and then take the average to obtain the identity-level feature denoted as CLIP-Memory. In this form, we replace the text encoder and achieve text-free CLIP-based person ReID. As shown in Fig. 1 (a), this is reasonable because the pre-trained CLIP text encoder and visual encoder are constrained by contrastive learning losses, where the outputs of the two encoders are aligned in a common feature space. What's more, the essence of the first-stage training of CLIP-ReID is to align text prompts with the outputs of the pre-trained visual encoder. Based on the above facts, we propose CLIP-Memory to address the lack of appropriate text descriptions when using pre-trained visual-language models. Considering that the CLIP-Memory is fixed during the training process, the resulting method will be only optimized for a specific set of training identities, while ignoring the appearance diversity of the same identity. Inspired by CoCoOp (Zhou et al. 2022a), we further design a Sequence-Specific Prompt (SSP) module, which updates this CLIP-Memory according to each video sequence when training.

On the other hand, we further propose a Temporal Memory Diffusion (TMD) module to capture temporal information. TMD is composed of Temporal Memory Construction (TMC) and Memory Diffusion (MD). Technically, it first takes the feature of each frame as input, and constructs a memory token. Then, based on memory tokens, TMC is employed to capture temporal information within the sequence. The updated memory token will store the context information of the video sequence. Next, MD combines the obtained temporal memory tokens with the original features, and diffuses the temporal memory to each token. Finally, we perform a Temporal Average Pooling (TAP) on the updated frame-level features to obtain robust sequence-level feature representations. Extensive experiments on three public ReID benchmarks demonstrate that our approach achieves promising results over most of previous methods.

In summary, our contributions are three folds:

- We propose a novel one-stage text-free CLIP-based learning framework named TF-CLIP for video-based person ReID. To our best knowledge, we are the first to extract identity-specific sequence features to replace the text features of CLIP. Meanwhile, we further design a Sequence-Specific Prompt (SSP) module to update the CLIP-Memory online.
- We propose a Temporal Memory Diffusion (TMD) module to capture temporal information. The frame-level memories in a sequence first communicate with each other to extract temporal information. The temporal in-

formation is then further diffused to each token, and finally aggregated to obtain more robust temporal features.

- Extensive experiments demonstrate that our proposed method shows superior performance over existing methods on three video-based person ReID datasets, *i.e.*, MARS, LS-VID and iLIDS-VID.

## Related Works

### Video-based Person ReID

Most of video-based person ReID methods exploit spatial and temporal cues for video representation learning. Early studies use RNNs/LSTMs (McLaughlin, Martinez del Rincon, and Miller 2016; Dai et al. 2018), 3D convolutions (Li, Zhang, and Huang 2019; Gu et al. 2020), temporal pooling (Zheng et al. 2016; Wu et al. 2018) and attention mechanisms (Liu et al. 2021c; Liu, Zhang, and Lu 2023), to extract temporal features. For example, Hou *et al.* (Hou et al. 2020) propose a temporal complementary learning network to extract complementary features of consecutive video frames. Different from these methods, we propose a Temporal Memory Diffusion (TMD) module to mine temporal memory and diffuse it to each frame.

In recent years, the vision Transformer has shown impressive results in person ReID compared with CNN-based methods. Liu *et al.* (Liu et al. 2021b) propose a trigeminal Transformer in spatial, temporal and spatial-temporal views to obtain more cues. He *et al.* (He et al. 2021) propose a hybrid framework based on dense interaction learning, which utilizes both CNNs and attention mechanisms to enhance multi-grained spatio-temporal modeling. Despite the impressive success of above approaches, they all employ a unimodal framework, where video encoders are trained to predict a fixed set of predefined labels. In contrast, we propose a new paradigm based on a visual-language multimodal learning framework for video-based person ReID.

### Visual-language Learning

Recently, visual-language joint learning methods (Jia et al. 2021; Yuan et al. 2021) (*e.g.*, CLIP) have demonstrated great potential in extracting generic visual representations. For example, Rao *et al.* (Rao et al. 2022) propose DenseCLIP for dense prediction by implicitly and explicitly leveraging the pre-trained knowledge from CLIP. Zhou *et al.* propose CoOp (Zhou et al. 2022b) and CoCoOp (Zhou et al. 2022a) for zero-shot image classification by prompt learning. Khatkhat *et al.* (Khattak et al. 2022) propose multi-modal prompt learning for both vision and language branches to improve alignments between vision and language representations. Li *et al.* (Li, Sun, and Li 2022) utilize CLIP for image-based person ReID with a two-stage training strategy. Different from the above methods, we design a one-stage training strategy to extend CLIP to the video-based person ReID task. On the other hand, several studies also try to extend the existing CLIP model to the video domain. Rasheed *et al.* (Rasheed et al. 2022) propose ViFi-CLIP to bridge the domain gap from images to videos. Wang *et al.* (Wang, Xing, and Liu 2021) propose ActionCLIP for action recognition by designing appropriate prompts. Ni *et al.* (Ni et al.

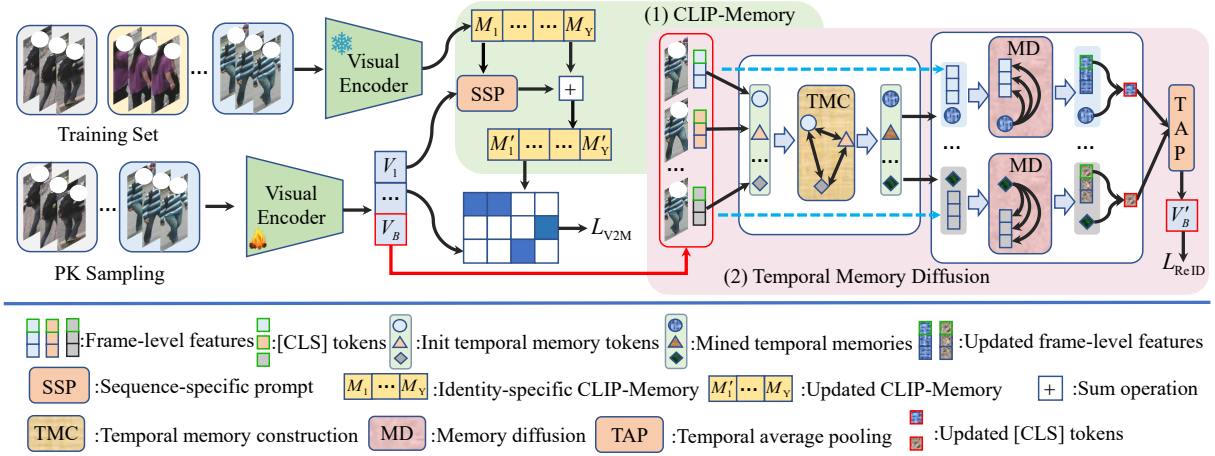


Figure 2: Illustration of the proposed TF-CLIP framework. It consists of two key modules, *i.e.*, a CLIP-Memory module and a Temporal Memory Diffusion (TMD) module. (1) In the CLIP-Memory module, we first extract the identity-specific sequence feature as CLIP-Memory to replace the text feature. Meanwhile, a Sequence-Specific Prompt (SSP) module is proposed to update the CLIP-Memory online. (2) In TMD, Temporal Memory Construction (TMC) is first proposed to capture temporal information based on the relations within the sequence. Then, Memory Diffusion (MD) further diffuses the temporal information to each token in the original features to obtain more robust sequence features.

2022) propose XCLIP and use a text prompt generation for better generalization. Inspired by the successful applications of CLIP, we further design a Temporal Memory Diffusion (TMD) module for video-based person ReID.

### Brief Reviews of CLIP and CLIP-ReID

CLIP consists of a visual encoder  $\mathcal{V}(\cdot)$  and a text encoder  $\mathcal{T}(\cdot)$ , jointly trained with contrastive learning to respectively map the input image and text into a unified representation space. Specifically, let  $\{img_i, text_i\}_{i=1}^B$  be a set of  $B$  training visual-language pairs within a batch, where  $img_i$  is an image and  $text_i$  is a corresponding text description. CLIP uses the above two encoders and combines two linear projecting layers to encode images and texts into corresponding image features and text features. Then, the similarity of the two features is calculated as:

$$S(img_i, text_i) = \mathcal{J}_v(\mathcal{V}(img_i)) \cdot \mathcal{J}_t(\mathcal{T}(text_i)), \quad (1)$$

where  $\mathcal{J}_v$  and  $\mathcal{J}_t$  are linear layers to project features into a joint feature space. Finally, two contrastive losses denoted as  $L_{v2t}$  and  $L_{t2v}$  are employed to train models (Radford et al. 2021). Thanks to  $L_{v2t}$  and  $L_{t2v}$ , the cross-modal features will be aligned. In the application of downstream tasks, the text usually takes the form of prompt, such as “A photo/video of a {class}.” where “{ }” is filled with a specific class name (*e.g.*, cat). Unfortunately, the labels in person ReID are integer indexes instead of specific texts.

To solve this problem, CLIP-ReID (Li, Sun, and Li 2022) introduces some identity-specific tokens to learn text descriptions. Specifically, the text prompt becomes “A photo of a  $[X]_1, [X]_2, \dots, [X]_{Num}$  person”. Then, in the additional training stage, CLIP-ReID fixes the parameters of  $\mathcal{V}(\cdot)$  and  $\mathcal{T}(\cdot)$ , and uses the  $L_{v2t}$  and  $L_{t2v}$  to optimize the learnable identity-specific tokens. In this way, CLIP-ReID learns

a corresponding ambiguous text description for each identity. Although CLIP-ReID achieves impressive results, it requires an additional training stage to learn text descriptions and does not consider the temporal information in videos.

### Our Method

To deal with the aforementioned problems, we propose a one-stage text-free framework named TF-CLIP. The proposed pipeline is depicted in Fig. 2, which consists of two main modules, a CLIP-Memory module to replace the original text branch and a Temporal Memory Diffusion (TMD) module to extend CLIP to video-based person ReID.

#### CLIP-Memory Module

Once the computation of the visual branch is completed, the next step of previous CLIP-based methods (Li, Sun, and Li 2022; Zhou et al. 2022b; Rao et al. 2022) is to generate the corresponding text features. Unlike them, we argue that learning text descriptions is not necessary for CLIP-based person ReID method. A CLIP-Memory can be used to replace the original text branch. Specifically, given a video-based person ReID training set  $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^{N_s}$  with labels  $y_i \in \{1, \dots, Y\}$ , the total number of the person tracklets is  $N_s$ . Taking identity  $y_i$  as an example, we first employ the pre-trained CLIP visual encoder for all video sequences with the identity  $y_i$  to extract identity-specific sequence features denoted as CLIP-Memory.

More specifically, given one of image sequences  $\{I_t\}_{t=1}^T$  containing  $T$  images belonging to the identity  $y_i$ , we split each frame  $I_t \in \mathbb{R}^{H \times W \times 3}$  into  $N_p$  non-overlapping patches  $\{I_{t,i}\}_{i=1}^{N_p} \in \mathbb{R}^{P^2 \times 3}$ , in which  $H$  and  $W$  represent the number of height and width, respectively.  $N_p = \frac{H}{P} \times \frac{W}{P}$  and  $P$  is the patch size. Then the patches are projected to token embeddings using a linear projection layer  $E_{emb} \in \mathbb{R}^{3P^2 \times D}$ ,

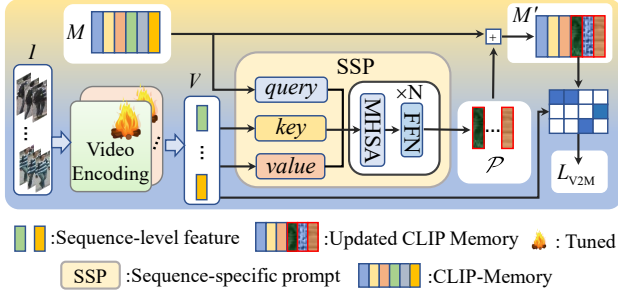


Figure 3: Illustration of our CLIP-Memory module.

where  $D$  represents the dimension of each token. Meanwhile, an extra learnable [CLS] token denoted as  $CLS_t \in \mathbb{R}^D$  is added to the embedded tokens for each frame. Thus, the input of the video encoder at the frame  $t$  is given by:

$$z_t^{Patch} = [CLS_t, E_{emb}^T I_{t,1}, \dots, E_{emb}^T I_{t,N_p}] + e^{sp}, \quad (2)$$

where  $e^{sp}$  is the spatial position embedding. Then the above embeddings are sequentially processed through CLIP. At last, the frame-level representation is defined as,

$$z_t = \mathcal{V}^*(z_t^{Patch}), \quad (3)$$

where the superscript  $*$  indicates that the parameters of this visual encoder  $\mathcal{V}$  are frozen during training. What's more, the obtained class token  $CLS_t$  is projected to a visual-language unified space via a Visual Project (VP) layer,  $f_t = VP(z_{t,0})$ , where  $f_t$  is the final representation of frame  $t$ . In practice,  $VP$  is implemented by a fully connected layer, represented as  $VP = W_{VP} \in \mathbb{R}^{D \times d}$ . Finally, a Temporal Average Pooling (TAP) is employed to obtain the sequence-level feature  $v_a$ . Once all the video sequence features belonging to the identity  $y_i$  are obtained, the average of them can represent the identity-specific feature  $M_{y_i}$  defined as CLIP-Memory,  $M_{y_i} = \frac{1}{N_i} \sum_{v_a \in y_i} v_a$ , where  $N_i$  is the number of sequences belonging to the identity  $y_i$ . When traversing the entire training set, we can get an identity-specific CLIP-memory  $M \in \mathbb{R}^{Y \times d}$ . Superior to previous text-associated CLIP-based methods, CLIP-Memory can exploit the image-text alignment potential of CLIP without using text information.

**Sequence-Specific Prompt.** In fact, if the CLIP-Memory is fixed during the training process, the framework will be only optimized for a specific set of training labels, while ignoring the appearance diversity of the same identity. To address the above problem, as shown in Fig. 3, we introduce a novel module: Sequence-Specific Prompt (SSP). The key idea is to generate a prompt conditioned on each input sequence to update the CLIP-Memory.

When training SSP, the PK sampling strategy (Hermans, Beyer, and Leibe 2017) is employed to form a mini-batch  $\{I^b\}_{b=1}^{B=P \times K}$ , which contains  $P$  different identities and  $K$  video sequences for each identity. After video encoding, the sequence-level features in the mini-batch can be expressed as  $V = [v_{y=1}^1, \dots, v_{y=1}^K, \dots, v_{y=P}^1, \dots, v_{y=P}^K]$ . Then, SSP takes the sequence-level feature  $V$  and CLIP-Memory  $M$  as inputs.

As shown in Fig. 3, the CLIP-Memory  $M$  is used as *query*, and the sequence-level feature  $V$  is used as *key* and *value*. Then, SSP employs the cross-attention mechanism to model the interactions between *query*, *key* and *value*, which can be achieved by:

$$\mathcal{P}_{y=1, \dots, y=P} = Transdecoder(query, key, value), \quad (4)$$

where  $\mathcal{P}$  represents the generated sequence-specific prompt. There are  $N$  blocks in *Transdecoder*, and each block consists of a Multi-Head Cross-Attention (MHCA) and a Feed-Forward Network (FFN). Subscripts  $y = 1, \dots, y = P$  represent the identities, which correspond to the index in the CLIP-Memory. We then update the identity-specific CLIP-Memory through a residual connection:

$$M'_{y=1, \dots, y=P} \leftarrow \mathcal{P}_{y=1, \dots, y=P} + M_{y=1, \dots, y=P}. \quad (5)$$

In this way, CLIP-Memory can be updated online. Finally, the updated CLIP-Memory  $M' \in \mathbb{R}^{Y \times d}$  and the sequence-level feature  $V \in \mathbb{R}^{PK \times d}$  are employed for contrastive learning. Compared with a fixed CLIP-Memory, SSP shifts the network's focus from a specific set of identities to each input sequence, thereby reducing overfitting and ultimately resulting in more discriminative features.

## Temporal Memory Diffusion

To capture temporal information, we further propose a Temporal Memory Diffusion (TMD) module. As shown in Fig. 4, TMD is mainly composed of Temporal Memory Construction (TMC) and Memory Diffusion (MD).

**Temporal Memory Construction.** Taking the last video sequence in the training batch as an example, which is defined as  $I^B = \{I_1^B, I_2^B, \dots, I_T^B\}$ , where  $B = P \times K$ . According to Eq. 2 and Eq. 3, the representation at the frame level denoted as  $z_t$  can be obtained through the video encoder. During the standard training process of CLIP, the [CLS] token  $z_{t,0}$  is used as the output of the visual encoder while the other tokens  $z_{t,others} = [z_{t,1}, \dots, z_{t,N_p}]$  are usually neglected. However, recent works (Zeng et al. 2022; Chen, Fan, and Panda 2021; He et al. 2022) have shown that  $z_{t,others}$  still retains enough semantic information and can be used as a feature map. Thus, Temporal Memory Construction (TMC) takes  $z_t \in \mathbb{R}^{(N_p+1) \times D}$  as input.

As shown in Fig. 4, we first construct a memory token  $s_t = \theta(\frac{1}{N_p+1} \sum_{i=0}^{N_p} z_{t,i})$  for each frame  $t$  based on frame-level features, where we get the initial memory token based on the average of all tokens in each frame, and  $\theta$  represents a liner projection. Then, to capture the temporal cues in the sequence and transfer the information of each frame into the memory tokens, a Multi-Head Self-Attention (MHSA) is computed over the initial memory tokens  $S = \{s_1, \dots, s_T\}$ :

$$S' = MHSA(LN(S)) + S, \quad (6)$$

where  $LN$  denotes Layer Normalization, and  $S' \in \mathbb{R}^{T \times D}$  are the temporal-aware memory tokens. TMC allows the frame-level memories to communicate with each other, and extracts temporal information based on the relations obtained from neighbors.

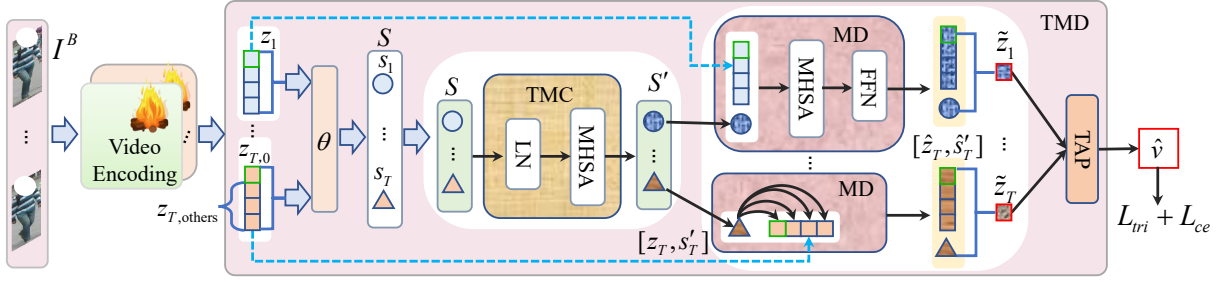


Figure 4: Illustration of the proposed Temporal Memory Diffusion (TMD) module.

**Memory Diffusion.** In order to make a better use of the temporal memory tokens, we design a diffusion-aggregation strategy named Memory Diffusion (MD). As shown in Fig. 4, the temporal memory token  $s'_t \in \mathbb{R}^{1 \times D}$  is first concatenated with each frame feature  $z_t \in \mathbb{R}^{(N_p+1) \times D}$  to become  $[z_t, s'_t] \in \mathbb{R}^{(N_p+2) \times D}$ , where  $[\cdot, \cdot]$  represents a concatenation of two features. Then, a MHSA followed by a FFN layer is employed to diffuse the temporal information to each token in the frame, resulting in  $[\hat{z}_t, \hat{s}'_t] \in \mathbb{R}^{(N_p+2) \times D}$ , which can be expressed as:

$$[\hat{z}_t, \hat{s}'_t] = FFN(MHSA([z_t, s'_t])) + [z_t, s'_t]. \quad (7)$$

Finally, all tokens in each frame are aggregated to obtain frame-level features  $\tilde{z}_t$  and a TAP is employed to further aggregate frame-level features into a sequence-level feature  $\hat{v}$ . The above processes can be formulated as:

$$\tilde{z}_t = \frac{1}{N_p + 2} \left( \sum_{i=0}^{N_p} \hat{z}_i + \hat{s}'_t \right), \quad (8)$$

$$\hat{v} = TAP([\tilde{z}_1, \dots, \tilde{z}_T]). \quad (9)$$

Superior to previous aggregation strategies, MD can diffuse the temporal memories to each token in the original features and obtain more robust sequence features.

### Optimization

In this work, we utilize the video-to-memory contrastive loss denoted by  $L_{V2M}$  to train the SSP module and the visual encoder in the CLIP-Memory. It is worth noting that the visual encoder used to compute CLIP-Memory does not update during training. Given the sequence-level feature  $V = [v_1, \dots, v_B]$  and the updated CLIP-Memory  $M'$ , it can be expressed as:

$$L_{V2M}^{(y_i)} = -\frac{1}{|P(y_i)|} \sum_{p \in P(y_i)} \log \frac{\exp(\langle v_p, M'_{y_i} \rangle)}{\sum_{j=1}^B \exp(\langle v_j, M'_{y_i} \rangle)}. \quad (10)$$

where  $P(y_i)$  represents all the positives for  $M'_{y_i}$  in the training batch and  $|\cdot|$  is its cardinality.  $\langle \cdot \rangle$  represents the cosine similarity function.  $B$  represents the number of video sequences in a batch.

What's more, the triplet loss (Hermans, Beyer, and Leibe 2017) denoted by  $L_{tri}$  and the label smooth cross-entropy

loss denoted by  $L_{ce}$  are employed to optimize the visual encoder and the TMD module. Finally, the losses used in our method are summarized as follows:

$$L_{total} = L_{V2M} + L_{tri} + L_{ce}. \quad (11)$$

## Experiments

### Datasets and Evaluation Protocols

We evaluate our proposed approach on three video-based person ReID benchmarks, including MARS (Zheng et al. 2016), LS-VID (Li et al. 2019) and iLIDS-VID (Wang et al. 2014). MARS is a large-scale benchmark which contains 20,715 videos with 1,261 identities. LS-VID is a new benchmark which collects 3,772 identities and includes 14,943 video tracklets captured by 15 cameras. iLIDS-VID is a small-scale dataset with 600 videos of 300 identities. In addition, we follow common practices and adopt the Cumulative Matching Characteristic (CMC) at Rank-k and mean Average Precision (mAP) to measure the performance.

### Experiment Settings

Our model is implemented on the PyTorch platform and trained with one NVIDIA Tesla A30 GPU (24G memory). The ViT-B/16 from CLIP (Radford et al. 2021) is used as the visual encoder's backbone. The number of blocks in SPP is set to  $N=2$ . During training, we sample 8 frames from each video sequence and each frame is resized to  $256 \times 128$ . In each mini-batch, we sample 4 identities, each with 4 tracklets. Thus, the number of images in a batch is  $4 \times 4 \times 8 = 128$ . We also adopt random flipping and random erasing (Zhong et al. 2020) for data augmentation. We train our framework for 60 epochs in total by the Adam optimizer (Kingma and Ba 2014). Following CLIP-ReID (Li, Sun, and Li 2022), we first warm up the model for 10 epochs with a linearly growing learning rate from  $5 \times 10^{-7}$  to  $5 \times 10^{-6}$ . Then, the learning rate is divided by 10 at the 30th and 50th epochs. The original sequence-level feature  $v \in \mathbb{R}^{1 \times 512}$  and the aggregated feature  $\hat{v} \in \mathbb{R}^{1 \times 768}$  are concatenated to obtain the final video representation during testing. The Euclidean distance is employed as the distance metric for ranking. The code is available at <https://github.com/AsuradaYuci/TF-CLIP>.

### Comparison with State-of-the-arts

In this section, we compare our methods with other methods on three video-based person ReID benchmarks. The results are shown in Tab. 1. These experimental results clearly

Methods	Source	Backbone	MARS		LS-VID		iLIDS-VID	
			mAP	Rank-1	mAP	Rank-1	Rank-1	Rank-5
STMP (Liu et al. 2019)	AAAI19	C	72.7	84.4	39.1	56.8	84.3	96.8
M3D (Li, Zhang, and Huang 2019)	AAAI19	C	74.1	84.4	40.1	57.7	74.0	94.3
GLTR (Li et al. 2019)	ICCV19	C	78.5	87.0	44.3	63.1	86.0	98.0
TCLNet (Hou et al. 2020)	ECCV20	C	85.1	89.8	70.3	81.5	86.6	-
MGH (Yan et al. 2020)	CVPR20	C	85.8	90.0	61.8	79.6	85.6	97.1
GRL (Liu et al. 2021c)	CVPR21	C	84.8	91.0	-	-	90.4	98.3
BiCnet-TKS (Hou et al. 2021)	CVPR21	C	86.0	90.2	75.1	84.6	-	-
CTL (Liu et al. 2021a)	CVPR21	C	86.7	91.4	-	-	89.7	97.0
STMN (Eom et al. 2021)	ICCV21	C	84.5	90.5	69.2	82.1	-	-
PSTA (Wang et al. 2021)	ICCV21	C	85.8	91.5	-	-	91.5	98.1
DIL (He et al. 2021)	ICCV21	T	87.0	90.8	-	-	92.0	98.0
STT (Zhang et al. 2021b)	Arxiv21	T	86.3	88.7	78.0	87.5	87.5	95.0
TMT (Liu et al. 2021b)	Arxiv21	T	85.8	91.2	-	-	91.3	98.6
CAVIT (Wu et al. 2022)	ECCV22	T	87.2	90.8	79.2	89.2	<u>93.3</u>	98.0
SINet (Bai et al. 2022)	CVPR22	C	86.2	91.0	79.6	87.4	92.5	-
MFA (Gu et al. 2022)	TIP22	C	85.0	90.4	78.9	88.2	<u>93.3</u>	<u>98.7</u>
DCCT (Liu et al. 2023)	TNNLS23	T	<u>87.5</u>	<u>92.3</u>	-	-	91.7	98.6
LSTRL (Liu, Zhang, and Lu 2023)	ICIG23	C	86.8	91.6	<u>82.4</u>	<u>89.8</u>	92.2	98.6
<b>TF-CLIP(Ours)</b>		T	<b>89.4</b>	<b>93.0</b>	<b>83.8</b>	<b>90.4</b>	<b>94.5</b>	<b>99.1</b>

Table 1: Comparison with typical CNN-based (C) and Transformer-based (T) methods on MARS, LS-VID and iLIDS-VID.

Model	Components			Params(M)	FLOPs(G)	MARS			LS-VID		
	CLIP-M	SSP	TMD			mAP	Rank-1	Rank-5	mAP	Rank-1	Rank-5
1	×	×	×	126.78	14.24	88.1	91.7	97.4	80.6	88.8	96.3
2	✓	×	×	86.94	11.26	88.4	91.6	97.9	80.3	87.7	95.7
3	✓	✓	×	94.21	12.11	88.8	92.1	97.6	81.3	89.9	96.2
4	✓	×	✓	97.02	11.72	89.2	92.3	97.7	81.6	88.9	96.4
5	×	×	✓	136.86	14.70	89.3	92.7	97.8	82.3	90.3	96.7
6	✓	✓	✓	104.26	12.53	89.4	93.0	97.9	83.8	90.4	97.1

Table 2: Comparison of different components and the computational cost on MARS and LS-VID.

confirm the superiority and effectiveness of the proposed method on large-scale ReID datasets.

**Results on MARS.** Tab. 1 shows that our proposed method achieves the best mAP of 89.4% and the best Rank-1 of 93.0% on MARS. We can observe that our method shows much better results than other methods. We note that most recent works (*e.g.*, DIL (He et al. 2021), CAVIT (Wu et al. 2022), etc.) are developed with Transformers. However, they are all based on unimodal pre-training. Different from them, we propose TF-CLIP based on multi-modal pre-training. Thus, our method surpasses CAVIT by 2.2% and 2.2% in terms of mAP and Rank-1 accuracy on MARS.

**Results on LS-VID.** It is observed that our method outperforms other state-of-the-arts on LS-VID. Among them, STMN (Eom et al. 2021) also uses a temporal memory, which attains 82.1% Rank-1 accuracy. Instead, we propose the TMD module to generate the temporal memory in the sequence. As a result, our method achieves 90.4% Rank-1 accuracy, which surpasses STMN by 8.3%.

**Results on iLIDS-VID.** As can be observed, Tab. 1 also demonstrates the significant superiority of the proposed model over existing methods on iLIDS-VID. Specially, our method achieves the best Rank-1 and Rank-5 accuracy of 94.2% and 99.1%, respectively. It surpasses the previous

best approaches MFA (Gu et al. 2022) and CAVIT (Wu et al. 2022) by 1.2% in terms of Rank-1 accuracy.

## Ablation Study

To verify the impact of each component in our model, we conduct several experiments on MARS and LS-VID, and show compared results in Tab. 2. “CLIP-M” stands for CLIP-Memory. *Model1* means that CLIP-ReID is directly applied to video-based person ReID as the baseline, in which a TAP is employed to obtain sequence-level features.

**Effectiveness of CLIP-Memory.** As can be seen from the first two rows in Tab. 2, *Model2* obtains a higher Rank-5 and mAP (by 0.5% and 0.3%) than *Model1* on MARS. Meanwhile, competitive results are achieved on LS-VID. The above results clearly demonstrate that it is feasible to use the CLIP-Memory to replace the text branch, which can inspire the subsequent extension of CLIP to other downstream tasks where text information is missing.

**Effectiveness of SSP.** As shown in Tab. 2, compared with *Model2*, *Model3* with SSP brings 0.5% mAP and 0.4% Rank-1 accuracy gains on MARS, respectively. What’s more, compared with *Model4*, *Model6* also improves the Rank-1 accuracy by 1.5% on LS-VID. These results clearly demonstrate that our proposed SSP can indeed improve the




Figure 5: Illustration of (a) the impact of different temporal fusion methods, (b) the impact of different layers in SSP, and (c) comparison w/o CLIP-Memory.

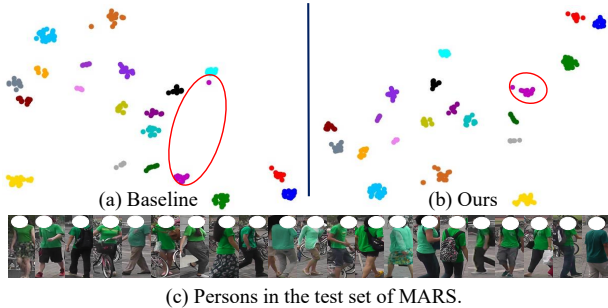


Figure 6: t-SNE visualization of the baseline and our method on the MARS test set. Different colored dots represent different identities. Best viewed in color.

performance. A reasonable explanation for this improvement is that updating the CLIP-Memory online helps the network to learn a more discriminative representation.

**Effectiveness of TMD.** As shown in Tab. 2, the proposed TMD improves the performance remarkably. Compared with *Model2*, *Model4* using TMD brings 0.8% mAP and 0.7% Rank-1 accuracy gains on MARS, respectively. What’s more, compared with *Model3*, *Model6* also brings 0.6% and 2.5% Rank-1 accuracy gains on MARS and LS-VID, respectively. The main reason is that our proposed TMD can capture temporal information in the sequence and obtain more robust sequence-level features.

**Comparison of different temporal fusion methods.** Following ActionCLIP (Wang, Xing, and Liu 2021), we further investigate the impact of different temporal fusion methods with *Model2* on MARS. The experimental results are shown in Fig. 5 (a). It can be observed that the proposed TMD is more suitable for video-based person ReID. For example, our method achieves 89.2% Rank-1 accuracy on MARS, which surpasses one-dimensional convolution (Conv1D) by 1.1%. The reason is that our proposed TMD not only extracts temporal information in the sequence, but also passes it to each token in the frame.

**The effect of different layers in SSP.** As shown in Fig. 5 (b), we also carry out experiments to investigate the effect of different layers in SSP with *Model3* on MARS. We can see that our method is not sensitive to this hyper-parameter. To balance the computation and performance, we finally choose  $N = 2$ , which achieves the best mAP accuracy of 88.8%.

**The necessity of CLIP-Memory.** In fact, we can train the proposed framework without using CLIP-Memory. The experimental results based on *Model2* and *Model4* are shown in Fig. 5 (c). It can be observed that not using CLIP-Memory will lead to a significant decrease in performance. For example, the unimodal baseline method achieves only 84.9% mAP accuracy on MARS, which is 3.5% lower than *Model2*. A plausible explanation for this performance drop is that the person ReID dataset is too small to adequately fine-tune the CLIP visual encoder when trained with unimodality. This reflects the superiority of multimodal training and the necessity of our CLIP-Memory.

## Visualization

To better understand the effect of the proposed TF-CLIP, some persons with similar appearances from MARS are chosen to visualize the feature distribution by t-SNE (Van der Maaten and Hinton 2008). As shown in Fig. 6 (c), these persons are wearing green clothes with small inter-person variations. Comparing Fig. 6 (a) and (b), we can find that the proposed TF-CLIP indeed helps to learn more discriminative embeddings, with which the intra-person variance is minimized and the intra-person variance is maximized.

## Conclusion

In this paper, we explore the application of vision-language pre-trained model to video-based person ReID without suitable textual descriptions. Specifically, we propose a novel one-stage text-free CLIP-based framework named TF-CLIP. To achieve text-free purpose, we propose CLIP-Memory to extract identity-specific sequence features to replace the text features. Meanwhile, we design a Sequence-Specific Prompt (SSP) to update the CLIP-Memory online. To capture temporal information, we further propose a Temporal Memory Diffusion (TMD) module. It first constructs frame-level memories and lets them communicate with each other in the sequence to extract the temporal information. Then the temporal memory is diffused into each token of the original frame, and aggregated to obtain more robust sequence features. Extensive experiments on three public ReID datasets demonstrate the effectiveness and superiority of our method.

## Acknowledgments

This work was supported in part by the National Key Research and Development Program of China (No. 2018AAA0102001), National Natural Science Foundation of China (No. 62101092) and Fundamental Research Funds for the Central Universities (No. DUT22QN228).

## References

- Bai, S.; Ma, B.; Chang, H.; Huang, R.; and Chen, X. 2022. Salient-to-broad transition for video person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7339–7348.
- Chen, C.-F. R.; Fan, Q.; and Panda, R. 2021. Crossvit: Cross-attention multi-scale vision transformer for image classification. In *Proceedings of the IEEE International Conference on Computer Vision*, 357–366.
- Chen, D.; Li, H.; Xiao, T.; Yi, S.; and Wang, X. 2018. 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*, 1169–1178.
- Dai, J.; Zhang, P.; Wang, D.; Lu, H.; and Wang, H. 2018. Video person re-identification by temporal residual learning. *IEEE Transactions on Image Processing*, 28(3): 1366–1377.
- Dai, J.; Zhang, P.; Wang, D.; Lu, H.; and Wang, H. 2019. Video Person Re-Identification by Temporal Residual Learning. *IEEE Transactions on Image Processing*, 28: 1366–1377.
- Eom, C.; Lee, G.; Lee, J.; and Ham, B. 2021. Video-based person re-identification with spatial and temporal memory networks. In *Proceedings of the IEEE International Conference on Computer Vision*, 12036–12045.
- Fu, Y.; Wang, X.; Wei, Y.; and Huang, T. 2019. Sta: Spatial-temporal attention for large-scale video-based person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 8287–8294.
- Gao, J.; and Nevatia, R. 2018. Revisiting temporal modeling for video-based person reid. *arXiv preprint arXiv:1805.02104*.
- Gu, X.; Chang, H.; Ma, B.; and Shan, S. 2022. Motion feature aggregation for video-based person re-identification. *IEEE Transactions on Image Processing*, 31: 3908–3919.
- Gu, X.; Chang, H.; Ma, B.; Zhang, H.; and Chen, X. 2020. Appearance-preserving 3d convolution for video-based person re-identification. In *Proceedings of the European Conference on Computer Vision*, 228–243.
- He, J.; Chen, J.-N.; Liu, S.; Kortylewski, A.; Yang, C.; Bai, Y.; and Wang, C. 2022. Transfg: A transformer architecture for fine-grained recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 852–860.
- He, T.; Jin, X.; Shen, X.; Huang, J.; Chen, Z.; and Hua, X.-S. 2021. Dense interaction learning for video-based person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, 1490–1501.
- Hermans, A.; Beyer, L.; and Leibe, B. 2017. In defense of the triplet loss for person re-identification. *arXiv:1703.07737*.
- Hou, R.; Chang, H.; Ma, B.; Huang, R.; and Shan, S. 2021. BiCnet-TKS: Learning efficient spatial-temporal representation for video person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014–2023.
- Hou, R.; Chang, H.; Ma, B.; Shan, S.; and Chen, X. 2020. Temporal complementary learning for video person re-identification. In *Proceedings of the European Conference on Computer Vision*, 388–405.
- Jia, C.; Yang, Y.; Xia, Y.; Chen, Y.-T.; Parekh, Z.; Pham, H.; Le, Q.; Sung, Y.-H.; Li, Z.; and Duerig, T. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, 4904–4916. PMLR.
- Ju, C.; Han, T.; Zheng, K.; Zhang, Y.; and Xie, W. 2022. Prompting visual-language models for efficient video understanding. In *Proceedings of the European Conference on Computer Vision*, 105–124. Springer.
- Khattak, M. U.; Rasheed, H.; Maaz, M.; Khan, S.; and Khan, F. S. 2022. Maple: Multi-modal prompt learning. *arXiv preprint arXiv:2210.03117*.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv:1412.6980*.
- Li, J.; Wang, J.; Tian, Q.; Gao, W.; and Zhang, S. 2019. Global-local temporal representations for video person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, 3958–3967.
- Li, J.; Zhang, S.; and Huang, T. 2019. Multi-scale 3d convolution network for video based person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 8618–8625.
- Li, S.; Sun, L.; and Li, Q. 2022. CLIP-ReID: Exploiting vision-language model for image re-identification without concrete text labels. *arXiv preprint arXiv:2211.13977*.
- Liu, J.; Zha, Z.-J.; Wu, W.; Zheng, K.; and Sun, Q. 2021a. Spatial-temporal correlation and topology learning for person re-identification in videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4370–4379.
- Liu, K.; Ma, B.; Zhang, W.; and Huang, R. 2015. A spatio-temporal appearance representation for video-based pedestrian re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, 3810–3818.
- Liu, X.; Yu, C.; Zhang, P.; and Lu, H. 2023. Deeply-coupled convolution-transformer with spatial-temporal complementary learning for video-based person re-identification. *arXiv preprint arXiv:2304.14122*.
- Liu, X.; Zhang, P.; and Lu, H. 2023. Video-based Person Re-identification with Long Short-Term Representation Learning. *arXiv preprint arXiv:2308.03703*.
- Liu, X.; Zhang, P.; Yu, C.; Lu, H.; Qian, X.; and Yang, X. 2021b. A video is worth three views: Trigeminal transformers for video-based person re-identification. *arXiv:2104.01745*.

- Liu, X.; Zhang, P.; Yu, C.; Lu, H.; and Yang, X. 2021c. Watching you: Global-guided reciprocal learning for video-based person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 13334–13343.
- Liu, Y.; Yuan, Z.; Zhou, W.; and Li, H. 2019. Spatial and temporal mutual promotion for video-based person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 8786–8793.
- McLaughlin, N.; Martinez del Rincon, J.; and Miller, P. 2016. Recurrent convolutional network for video-based person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1325–1334.
- Ni, B.; Peng, H.; Chen, M.; Zhang, S.; Meng, G.; Fu, J.; Xiang, S.; and Ling, H. 2022. Expanding language-image pretrained models for general video recognition. In *Proceedings of the European Conference on Computer Vision*, 1–18. Springer.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 8748–8763. PMLR.
- Rao, Y.; Zhao, W.; Chen, G.; Tang, Y.; Zhu, Z.; Huang, G.; Zhou, J.; and Lu, J. 2022. Denseclip: Language-guided dense prediction with context-aware prompting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 18082–18091.
- Rasheed, H.; Khattak, M. U.; Maaz, M.; Khan, S.; and Khan, F. S. 2022. Fine-tuned CLIP models are efficient video learners. *arXiv preprint arXiv:2212.03640*.
- Subramaniam, A.; Nambiar, A.; and Mittal, A. 2019. Co-segmentation inspired attention networks for video-based person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, 562–572.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9: 2579–2605.
- Wang, M.; Xing, J.; and Liu, Y. 2021. Actionclip: A new paradigm for video action recognition. *arXiv preprint arXiv:2109.08472*.
- Wang, T.; Gong, S.; Zhu, X.; and Wang, S. 2014. Person Re-identification by video ranking. In *Proceedings of the European Conference on Computer Vision*, 688–703.
- Wang, Y.; Zhang, P.; Gao, S.; Geng, X.; Lu, H.; and Wang, D. 2021. Pyramid spatial-temporal aggregation for video-based person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 12026–12035.
- Wu, J.; He, L.; Liu, W.; Yang, Y.; Lei, Z.; Mei, T.; and Li, S. Z. 2022. CAViT: Contextual alignment vision transformer for video object re-identification. In *Proceedings of the European Conference on Computer Vision*, 549–566. Springer.
- Wu, Y.; Lin, Y.; Dong, X.; Yan, Y.; Ouyang, W.; and Yang, Y. 2018. 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*, 5177–5186.
- Xu, H.; Ghosh, G.; Huang, P.-Y.; Okhonko, D.; Aghajanyan, A.; Metz, F.; Zettlemoyer, L.; and Feichtenhofer, C. 2021. Videoclip: Contrastive pre-training for zero-shot video-text understanding. *arXiv preprint arXiv:2109.14084*.
- Xu, S.; Cheng, Y.; Gu, K.; Yang, Y.; Chang, S.; and Zhou, P. 2017. Jointly attentive spatial-temporal pooling networks for video-based person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, 4733–4742.
- Yan, Y.; Qin, J.; Chen, J.; Liu, L.; Zhu, F.; Tai, Y.; and Shao, L. 2020. Learning multi-granular hypergraphs for video-based person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2899–2908.
- Yuan, L.; Chen, D.; Chen, Y.-L.; Codella, N.; Dai, X.; Gao, J.; Hu, H.; Huang, X.; Li, B.; Li, C.; et al. 2021. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*.
- Zeng, W.; Jin, S.; Liu, W.; Qian, C.; Luo, P.; Ouyang, W.; and Wang, X. 2022. Not all tokens are equal: Human-centric visual analysis via token clustering transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11101–11111.
- Zhang, G.; Zhang, P.; Qi, J.; and Lu, H. 2021a. Hat: Hierarchical aggregation transformers for person re-identification. In *Proceedings of the 29th ACM International Conference on Multimedia*, 516–525.
- Zhang, T.; Wei, L.; Xie, L.; Zhuang, Z.; Zhang, Y.; Li, B.; and Tian, Q. 2021b. Spatiotemporal transformer for video-based person re-identification. *arXiv:2103.16469*.
- Zhang, Z.; Lan, C.; Zeng, W.; and Chen, Z. 2020. Multi-granularity reference-aided attentive feature aggregation for video-based person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 10407–10416.
- Zheng, L.; Bie, Z.; Sun, Y.; Wang, J.; Su, C.; Wang, S.; and Tian, Q. 2016. Mars: A video benchmark for large-scale person re-identification. In *Proceedings of the European Conference on Computer Vision*, 868–884.
- 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*, 13001–13008.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022a. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 16816–16825.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022b. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9): 2337–2348.
- Zhu, J.; Lai, S.; Chen, X.; Wang, D.; and Lu, H. 2023. Visual prompt multi-modal tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9516–9526.