

# MambaPro: Multi-Modal Object Re-Identification with Mamba Aggregation and Synergistic Prompt

Yuhao Wang<sup>1</sup>, Xuehu Liu<sup>2</sup>, Tianyu Yan<sup>1</sup>, Yang Liu<sup>1,4</sup>, Aihua Zheng<sup>3,4</sup>,  
Pingping Zhang<sup>1,4\*</sup>, Huchuan Lu<sup>1</sup>

<sup>1</sup>School of Future Technology, School of Artificial Intelligence, Dalian University of Technology

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

<sup>3</sup>School of Artificial Intelligence, Anhui University

<sup>4</sup>Anhui Provincial Key Laboratory of Multimodal Cognitive Computation, Anhui University

{924973292, 2981431354}@mail.dlut.edu.cn, liuxuehu@whut.edu.cn,

{ly, zhpp, lhchuan}@dlut.edu.cn, ahzheng214@foxmail.com

## Abstract

Multi-modal object Re-Identification (ReID) aims to retrieve specific objects by utilizing complementary image information from different modalities. Recently, large-scale pre-trained models like CLIP have demonstrated impressive performance in traditional single-modal object ReID tasks. However, they remain unexplored for multi-modal object ReID. Furthermore, current multi-modal aggregation methods have obvious limitations in dealing with long sequences from different modalities. To address above issues, we introduce a novel framework called MambaPro for multi-modal object ReID. To be specific, we first employ a Parallel Feed-Forward Adapter (PFA) for adapting CLIP to multi-modal object ReID. Then, we propose the Synergistic Residual Prompt (SRP) to guide the joint learning of multi-modal features. Finally, leveraging Mamba’s superior scalability for long sequences, we introduce Mamba Aggregation (MA) to efficiently model interactions between different modalities. As a result, MambaPro could extract more robust features with lower complexity. Extensive experiments on three multi-modal object ReID benchmarks (i.e., RGBNT201, RGBNT100 and MSVR310) validate the effectiveness of our proposed methods.

## Introduction

Object Re-Identification (ReID) aims to re-identify specific objects across non-overlapping cameras. Due to its wide applications, object ReID has advanced significantly in recent years (Liu et al. 2023; Wang et al. 2024a; Liu et al. 2024a, 2021; Zhang et al. 2021; Yu et al. 2024a). However, single-modal object ReID has many limitations in challenging scenarios (Li et al. 2020; Zheng et al. 2021), such as lighting changes, shadows and low image resolutions. Under such extreme conditions, single-modal object ReID methods may extract misleading features (Li et al. 2020), resulting in the loss of discriminative information (Wang et al. 2023). Fortunately, multi-modal object ReID has demonstrated a promising capability in addressing these challenges (Wang et al. 2023; Shi et al. 2024a,b). With complementary im-

\*Corresponding author (zhpp@dlut.edu.cn).

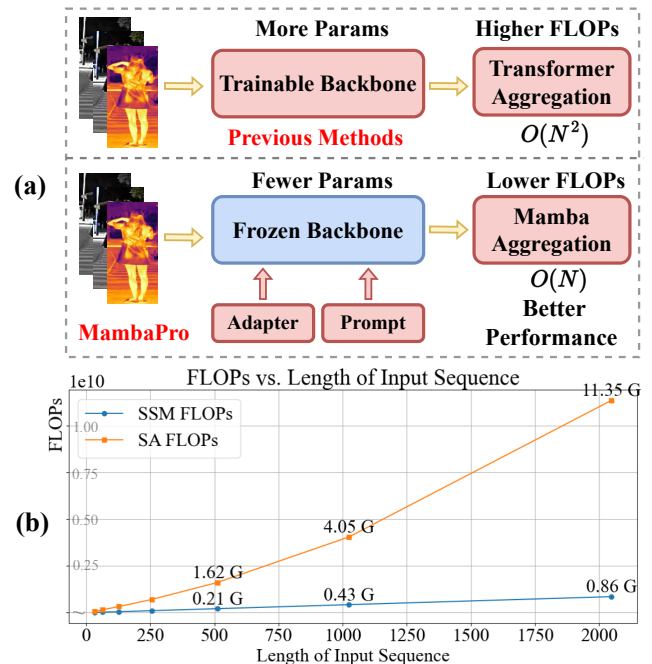


Figure 1: (a) Comparison between previous methods and MambaPro. (b) FLOPs comparison between SSM and SA.

age information from different modalities, multi-modal object ReID methods can obtain more robust feature representations (Zheng et al. 2023; Wang et al. 2023; Crawford et al. 2023; Yang, Chen, and Ye 2023, 2024). Meanwhile, with the widespread application of Transformers (Vaswani et al. 2017), the framework for multi-modal object ReID also changes from CNN-based methods (Wang et al. 2022; Zheng et al. 2023) to Transformer-based methods (Wang et al. 2023; Zhang et al. 2024). Although these methods show promising performance, they still face challenges in computational complexity as shown in the top of Fig. 1 (a).

Recently, large-scale pre-trained models like CLIP (Radford et al. 2021) have demonstrated strong generalization capabilities in downstream tasks, such as single-modal object ReID (Li, Sun, and Li 2023). However, full fine-tuning

of CLIP leads to the high computational complexity and catastrophic forgetting (French 1999). To address above issues, Parameter-Efficient Fine-Tuning (PEFT) methods like adapter and prompt tuning show competitive performance with fewer parameters and FLOPs. Meanwhile, current multi-modal aggregation methods face challenges in managing the high computational complexity that arises from interactions between different modalities. As shown in Fig. 1 (b), Self-Attention (SA)-based aggregation methods (Wang et al. 2023; Yin et al. 2023; Pan et al. 2023) struggle with quadratic complexity when handling long sequences from multiple modalities. Fortunately, the Selective State Space Models (SSM) in Mamba (Gu and Dao 2023) offer an efficient solution. They provide a superior scalability for long sequences by reducing complexity to linear levels (Gu, Goel, and Ré 2021). Motivated by the above observations, we propose MambaPro, a novel framework that combines CLIP-driven synergistic prompt tuning with Mamba aggregation.

To be specific, our MambaPro consists of three key components: Parallel Feed-Forward Adapter (PFA), Synergistic Residual Prompt (SRP) and Mamba Aggregation (MA). Technically, we start by inserting the PFA into the frozen image encoder of CLIP to facilitate the transfer of pre-trained knowledge into the ReID task. To achieve this, the PFA is implemented as a parallel branch to the Feed-Forward Network (FFN). This parallel design not only maintains the original multi-modal feature representations but also supports flexible knowledge transfer across modalities. By combining features derived from the CLIP pre-trained model with refinements from the adapter, the framework can extract more robust representations for each modality. Then, we introduce the SRP to guide the joint learning of multi-modal features. To be specific, we propose a Synergistic Prompt (SP) to achieve the synergistic transfer with modality-specific prompts, facilitating effective exchanges of discriminative multi-modal information. Furthermore, we introduce a Residual Prompt (RP) to layer-wisely aggregate multi-modal information with residual refinements. As a result, our SRP can significantly enhance the information synergism. Finally, we propose the MA to efficiently model interactions between long sequences from different modalities. Specifically, the stacked MA blocks integrate complementary features from two aspects: intra-modality and inter-modality. With the cooperation of these two aspects, our MA can fully capture discriminative information with linear complexity. As a result, our framework can extract more robust multi-modal features with lower complexity. Experiments on three multi-modal object ReID benchmarks clearly demonstrate the effectiveness of our proposed methods.

In summary, our contributions are as follows:

- We propose MambaPro, a novel framework for multi-modal object ReID. To our best knowledge, we are the first to introduce CLIP into multi-modal object ReID with Mamba aggregation and synergistic prompt tuning.
- We develop a Synergistic Residual Prompt (SRP) to guide the joint learning of multi-modal features, which effectively facilitate knowledge transfer and modality interactions with fewer parameters and FLOPs.

- We introduce a Mamba Aggregation (MA) to fully integrate the complementary information within and across different modalities with linear complexity.
- Extensive experiments are performed on three multi-modal object ReID benchmarks. The results fully validate the effectiveness of our proposed methods.

## Related Work

### Multi-Modal Object ReID

Traditional object ReID aims to extract features with single-modal input, such as RGB images or depth maps. In the past few years, single-modal object ReID has achieved remarkable progress. However, relying solely on single-modal input, these methods may extract limited information in complex scenarios. Meanwhile, the generalization capability of multi-modal data has sparked increasing attention in multi-modal object ReID (Zheng et al. 2021; Li et al. 2020; Yin et al. 2023; He et al. 2023; Wang et al. 2023). However, multi-modal input introduces additional challenges, such as distribution gaps among different modalities (Liang et al. 2022) and modality laziness (Crawford et al. 2023). To address above challenges, Zheng *et al.* (Zheng et al. 2021) propose to learn robust features through a progressive fusion for multi-modal person ReID. For multi-modal vehicle ReID, Li *et al.* (Li et al. 2020) propose to learn balanced multi-modal features. Based on Transformers (Vaswani et al. 2017), Pan *et al.* (Pan et al. 2023) employ a feature hybrid mechanism to control modal-specific information. Recently, Zhang *et al.* (Zhang et al. 2024) propose to select diverse tokens with the cooperation of multi-modal information. Despite promising results, these methods still face challenges in generalization capability and model complexity. Especially for previous Transformer-based methods, full fine-tuning introduces high computational complexity. Thus, with the strong transfer capabilities of large-scale pre-trained models, we introduce the powerful CLIP into multi-modal object ReID with fewer trainable parameters and FLOPs.

### Parameter-Efficient Fine-Tuning

With the development of foundational models (Yang et al. 2024), the transfer of pre-trained knowledge is attracting increasing attention (Diao et al. 2024a,b). Currently, mainstream PEFT methods can be divided into three categories: prompt tuning, adapter tuning and LoRA tuning. Prompt tuning (Lester, Al-Rfou, and Constant 2021) usually inserts layer-wise learnable tokens into the frozen backbone and guides the model to learn task-specific knowledge. Adapter tuning (Houlsby et al. 2019) introduces a plug-and-play module into the backbone, typically consisting of an MLP or attention module. LoRA tuning (Hu et al. 2021) uses a low-rank approximation as a side branch to reshape the backbone. With the above methods, PEFT demonstrates promising performance with limited resources. However, knowledge transfer for multi-modal fusion has been insufficiently explored, particularly in multi-modal object ReID. Therefore, we propose the PFA and SRP to leverage adapter tuning and prompt tuning for multi-modal object ReID. With the parallel adapters and synergistic transformations across

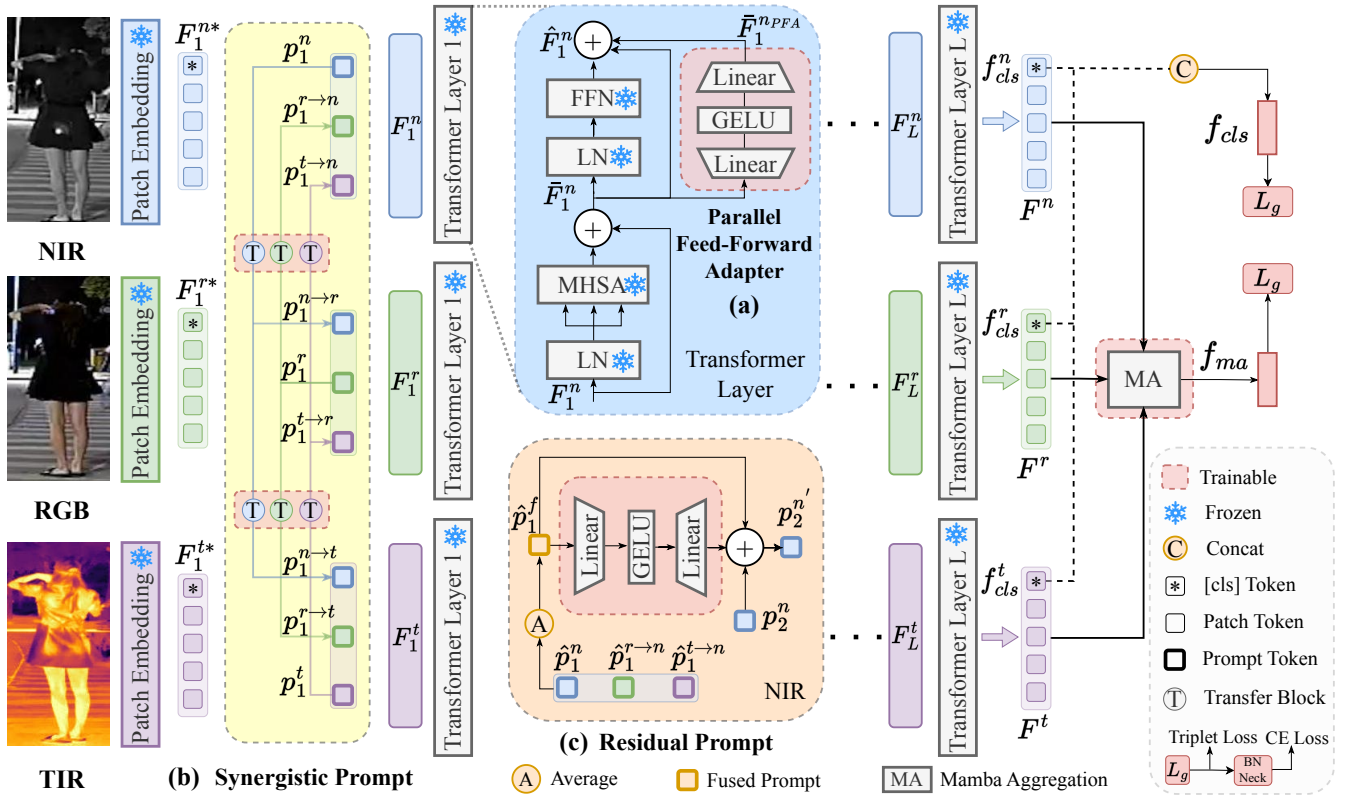


Figure 2: The overall framework of MambaPro. The Parallel Feed-Forward Adapter (PFA) is first introduced to transfer pre-trained knowledge into the ReID task. Then, the Synergistic Residual Prompt (SRP) is inserted to guide the progressive fusion of multi-modal features. Finally, the Mamba Aggregation (MA) is proposed to model interactions of long sequences from different modalities. With the proposed modules, our framework can obtain more robust features with low computational complexity.

multi-modal prompts, we can fully integrate the complementary information with lower complexity.

## Visual State Space Models

Recently, Mamba (Gu and Dao 2023) draws considerable attention (Wan et al. 2024; Xu et al. 2024) with the superior scalability of State Space Models (SSM) (Gu et al. 2021). By making the Structured State Space Sequence Model (S4) (Gu, Goel, and Ré 2021) data-dependent with a selection mechanism, Mamba outperforms Transformers in handling long sequences with linear complexity. In computer vision, ViS4mer (Islam and Bertasius 2022) and TranS4mer (Islam et al. 2023) demonstrate the effectiveness of SSM. To fully model the interactions between image patches, Vision Mamba (Zhu et al. 2024) and VMamba (Liu et al. 2024b) employ different scanning strategies. However, most of these methods focus on single-modal tasks. Additionally, attention-based aggregation methods often face quadratic complexity. To address these issues, we design the intra-modality and inter-modality aggregation with Mamba, which together form the complete MA block. This approach integrates complementary information with linear complexity and complements the progressive fusion from PFA and SRP, enhancing feature robustness for ReID.

## Method

As shown in Fig. 2, our framework comprises three main components: Parallel Feed-Forward Adapter (PFA), Synergistic Residual Prompt (SRP) and Mamba Aggregation (MA). With the frozen image encoder of CLIP as the shared backbone, we extract discriminative features from different modalities, i.e., RGB, Near Infrared (NIR) and Thermal Infrared (TIR). Details of the each component are as follows.

### Parallel Feed-Forward Adapter

To transfer CLIP’s knowledge to ReID tasks, we introduce the Parallel Feed-Forward Adapter (PFA). Previous methods (Diao et al. 2024b) typically insert adapters sequentially into networks. However, this can disrupt the original information flow, leading to suboptimal feature transformations. Especially in multi-modal settings, it is crucial to maintain the integrity of input features. To address these issues, our PFA is designed as a parallel branch to the Feed-Forward Network (FFN) rather than as a sequential operation. Taking the NIR modality as an example, as shown in Fig. 2 (a), the tokenized features  $F_1^n$  are first processed through the initial Transformer layer with the following equation:

$$\bar{F}_1^n = \mathcal{M}(\mathcal{L}(F_1^n)) + F_1^n, \quad (1)$$

where  $\mathcal{L}$  represents the Layer Normalization (LN) (Ba, Kiros, and Hinton 2016) and  $\mathcal{M}$  denotes the multi-head self-attention (Dosovitskiy et al. 2020). The output  $F_1^n$  is then fed into the PFA, undergoing the following transformation:

$$\bar{F}_1^{nPFA} = \xi(\delta(\xi(\bar{F}_1^n))). \quad (2)$$

Here,  $\xi$  is a linear layer and  $\delta$  is the GELU activation function (Hendrycks and Gimpel 2016). Finally, we integrate the output  $\bar{F}_1^{nPFA}$  with the FFN, resulting in the feature  $\hat{F}_1^n$ :

$$\hat{F}_1^n = \mathcal{F}(\mathcal{L}(\bar{F}_1^n)) + \bar{F}_1^{nPFA} + \bar{F}_1^n, \quad (3)$$

where  $\mathcal{F}$  represents the FFN. This parallel structure preserves the original feature representations while enabling flexible and efficient knowledge transfer across different modalities. Additionally, PFA employs an ascending-then-descending pattern in linear layers. This design enhances its ability to adapt to complex multi-modal feature distributions, supporting more stable knowledge transfer.

### Synergistic Residual Prompt

To promote the joint learning of multi-modal features, we introduce the Synergistic Residual Prompt (SRP). Different from previous single modal prompt-based methods, our SRP is designed to integrate the modality-specific knowledge with synergistic transformations. Besides, current multi-modal prompt-based methods (Li et al. 2023; Khattak et al. 2023) often discard prompts from earlier layers, potentially resulting in the loss of discriminative information. To address these issues, our SRP introduces two key components: Synergistic Prompt (SP) and Residual Prompt (RP). The SP transfers modality-specific knowledge across different modalities, while the RP aggregates multi-modal information across layers with residual refinements.

**Synergistic Prompt.** Without loss of generality, we take the NIR modality as an example. In Fig. 2 (b), the NIR image is first tokenized by the patch embedding (Vaswani et al. 2017). The patch tokens  $f_{pa}^n \in \mathbb{R}^{D \times N_{pa}}$  are concatenated with the class token  $f_{cls}^n \in \mathbb{R}^D$  to form  $F_1^{n*} \in \mathbb{R}^{D \times (1+N_{pa})}$ :

$$F_1^{n*} = [f_{cls}^n, f_{pa}^n], \quad (4)$$

where  $[\cdot]$  is the concatenation operation. For the prompt tokens in NIR modality, they are composed of three parts:  $p_1^n, p_1^{r \rightarrow n}, p_1^{t \rightarrow n} \in \mathbb{R}^{D \times N_{pr}}$ . Here,  $p_1^n$  is the randomly initialized NIR-specific prompt, while  $p_1^{r \rightarrow n}$  and  $p_1^{t \rightarrow n}$  are transferred from RGB and TIR modalities, respectively. Finally, all of them will be concatenated to form the input  $F_1^n \in \mathbb{R}^{D \times (1+N_{pa}+3N_{pr})}$  of the first layer as follows:

$$F_1^n = [F_1^{n*}, p_1^n, p_1^{r \rightarrow n}, p_1^{t \rightarrow n}]. \quad (5)$$

Here,  $D$  is the embedding dimension and  $N_{pa}$  is the number of patches.  $N_{pr}$  is the number of prompts in each modality.

$$p_1^{x \rightarrow n} = \mathcal{T}(p_1^x) = \xi(\delta(\xi(p_1^x))), \quad (6)$$

where  $x$  can be  $r$  or  $t$ .  $\mathcal{T}$  is the transfer block composed of linear layers and a GELU activation function. Then, the output  $\hat{F}_1^n$  of the first Transformer layer  $\Phi_1$  is calculated with the input  $F_1^n$  with the following equation:

$$\hat{F}_1^n = \Phi_1(F_1^n). \quad (7)$$

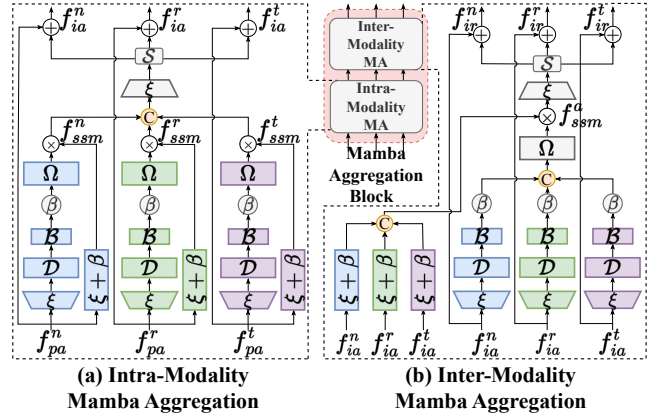


Figure 3: Details of our proposed Mamba Aggregation.

Other modalities undergo similar operations. As a result, we can leverage SP to facilitate the transfer of discriminative modality-specific information across different modalities.

**Residual Prompt.** As shown in Fig. 2 (c), we first extract prompt tokens  $\hat{p}_1^n, \hat{p}_1^{r \rightarrow n}, \hat{p}_1^{t \rightarrow n}$  from  $\hat{F}_1^n$ , with NIR modality as an example. Then, we average these prompts to form the fused prompts  $\hat{p}_1^f \in \mathbb{R}^{D \times N_{pr}}$ . Afterwards, we add the fused prompts  $\hat{p}_1^f$  with linear layers to the new NIR-specific prompt  $p_2^n$  to form the final prompts for the second layer:

$$p_2^n = p_2^n + \xi(\delta(\xi(\hat{p}_1^f))), \hat{p}_1^f = \frac{\hat{p}_1^n + \hat{p}_1^{r \rightarrow n} + \hat{p}_1^{t \rightarrow n}}{3}. \quad (8)$$

Finally, the input  $F_2^n$  of the second layer  $\Phi_2$  is calculated as:

$$F_2^n = [F_2^{n*}, p_2^n, \mathcal{T}(p_2^r), \mathcal{T}(p_2^t)]. \quad (9)$$

Other modalities and layers are processed in the same way:

$$\hat{F}_l^n = \Phi_l(F_l^n), F_l^n = [F_l^{n*}, p_l^n, \mathcal{T}(p_l^r), \mathcal{T}(p_l^t)], \quad (10)$$

$$\hat{F}_l^r = \Phi_l(F_l^r), F_l^r = [F_l^{r*}, \mathcal{T}(p_l^n), p_l^r, \mathcal{T}(p_l^t)], \quad (11)$$

$$\hat{F}_l^t = \Phi_l(F_l^t), F_l^t = [F_l^{t*}, \mathcal{T}(p_l^n), \mathcal{T}(p_l^r), p_l^t], \quad (12)$$

where  $l \in \{2, 3, \dots, L\}$ . With the synergistic transfer of modality-specific knowledge across different modalities and the residual refinements of multi-modal information, our SRP fully enhances the information synergism, promoting the feature discrimination. After the feature extraction from the backbone, we obtain  $F^n, F^r, F^t \in \mathbb{R}^{D \times (1+N_{pa})}$  for each modality. Then, we extract class tokens  $f_{cls}^n, f_{cls}^r, f_{cls}^t$  and concatenate them to form  $f_{cls} \in \mathbb{R}^{3D}$  for supervision.

### Mamba Aggregation

To efficiently model interactions between tokenized sequences from different modalities, we introduce Mamba Aggregation (MA) as a fusion strategy. As depicted in Fig. 3, MA is essentially composed of two parts: intra-modality MA and inter-modality MA. The intra-modality MA is designed to capture the discriminative information within each

modality. While the inter-modality MA is introduced to fully integrate the complementary features across different modalities. In the intra-modality MA, we first extract patch tokens  $f_{pa}^n, f_{pa}^r, f_{pa}^t \in \mathbb{R}^{D \times N_{pa}}$  from  $F^n, F^r$  and  $F^t$ , respectively. In the left part of Fig. 3, each of them will undergo the Mamba block composed of Depth-Wise Convolution (DWConv) (Chollet 2017)  $\mathcal{D}$ , Batch Normalization (BN) (Ioffe and Szegedy 2015)  $\mathcal{B}$ , SSM  $\Omega$  and SiLU activation function (Elfwing, Uchibe, and Doya 2018)  $\beta$  with the following transformations:

$$\Theta(\mathcal{X}) = \beta(\mathcal{B}(\mathcal{D}(\xi(\mathcal{X})))), \Psi(\mathcal{X}) = \beta(\xi(\mathcal{X})), \quad (13)$$

$$f_{ssm}^M = \Omega(\Theta(f_{pa}^M)) \times \Psi(f_{pa}^M), M \in \{n, r, t\}. \quad (14)$$

Here,  $M$  denotes the modality. Then, features from different Mamba blocks are concatenated along the token dimension. After passing through a linear layer, the aggregated features are split by modality for the intra-modality MA:

$$f_{ia}^n, f_{ia}^r, f_{ia}^t = \mathcal{S}(\xi([f_{ssm}^n, f_{ssm}^r, f_{ssm}^t]) + [f_{pa}^n, f_{pa}^r, f_{pa}^t]), \quad (15)$$

where  $f_{ia}^n, f_{ia}^r, f_{ia}^t \in \mathbb{R}^{D \times N_{pa}}$ . Here,  $\mathcal{S}$  denotes the feature splitting along the token dimension. Thus, through a separable SSM for each modality, the intra-modality MA effectively captures interactions within each modality.

Furthermore, as shown in the right part of Fig. 3,  $f_{ia}^n, f_{ia}^r$  and  $f_{ia}^t$  will be sent to the inter-modality MA as follows:

$$f_{ssm}^a = \Omega([\Theta(f_{ia}^n), \Theta(f_{ia}^r), \Theta(f_{ia}^t)]) \times [\Psi(f_{ia}^n), \Psi(f_{ia}^r), \Psi(f_{ia}^t)]). \quad (16)$$

Here,  $f_{ssm}^a \in \mathbb{R}^{D \times 3N_{pa}}$  represents the aggregated features obtained by modeling all modality tokens as the long sequences. Finally, we split the aggregated features into different modalities for the next MA block:

$$f_{ir}^n, f_{ir}^r, f_{ir}^t = \mathcal{S}(\xi(f_{ssm}^a) + [f_{ia}^n, f_{ia}^r, f_{ia}^t]). \quad (17)$$

Through the cooperation of the intra-modality MA and inter-modality MA, MA fully models the interactions of patch tokens within and across modalities. After the stacked MA blocks, we concatenate the class tokens and the average of patch tokens of each modality. Then, we use a LN to stabilize the feature learning. With a linear reduction, we obtain the final features  $f_{final}^M, M \in \{n, r, t\}$  for each modality:

$$f_{final}^M = \xi(\mathcal{L}([f_{cls}^M, \mathcal{A}(f_{ir}^M)])). \quad (18)$$

Here,  $\mathcal{A}$  represents the average operation. Finally, we concatenate the features from different modalities to form  $f_{ma} \in \mathbb{R}^{3D}$  for the loss supervision:

$$f_{ma} = [f_{final}^n, f_{final}^r, f_{final}^t]. \quad (19)$$

With the above aggregation strategy, we can fully model the interactions across modalities with linear complexity.

## Objective Functions

As illustrated in Fig. 2, our objective function comprises two parts: losses for the image encoder of CLIP and the MA. For the backbone and MA, they are both supervised by the label smoothing cross-entropy loss (Szegedy et al. 2016) and triplet loss (Hermans, Beyrer, and Leibe 2017):

$$\mathcal{L}_g = \lambda_1 \mathcal{L}_{ce} + \lambda_2 \mathcal{L}_{tri}. \quad (20)$$

Finally, the total loss for our framework can be defined as:

$$\mathcal{L}_{total} = \mathcal{L}_g^{CLIP} + \mathcal{L}_g^{MA}. \quad (21)$$

Methods		RGBNT201			
		mAP	R-1	R-5	R-10
Single	MUDeep (Qian et al. 2017)	23.8	19.7	33.1	44.3
	PCB (Sun et al. 2018)	32.8	28.1	37.4	46.9
	OSNet (Zhou et al. 2019)	25.4	22.3	35.1	44.7
	CAL (Rao et al. 2021)	27.6	24.3	36.5	45.7
Multi	HAMNet (Li et al. 2020)	27.7	26.3	41.5	51.7
	PFNet (Zheng et al. 2021)	38.5	38.9	52.0	58.4
	IEEE (Wang et al. 2022)	49.5	48.4	59.1	65.6
	DENet (Zheng et al. 2023)	42.4	42.2	55.3	64.5
	LRMM (Wu et al. 2025)	52.3	53.4	64.6	73.2
	UniCat* (Crawford et al. 2023)	57.0	55.7	-	-
	HTT* (Wang et al. 2024b)	71.1	73.4	83.1	87.3
	TOP-ReID* (Wang et al. 2023)	<u>72.3</u>	<u>76.6</u>	<u>84.7</u>	<u>89.4</u>
	EDITOR* (Zhang et al. 2024)	66.5	68.3	81.1	88.2
	RSCNet* (Yu et al. 2024b)	68.2	72.5	-	-
MambaPro $\ddagger$		<b>78.9</b>	<b>83.4</b>	<b>89.8</b>	<b>91.9</b>

Table 1: Performance comparison on RGBNT201. The best results are highlighted in bold, with the second-best underlined. Symbols:  $\ddagger$  indicates CLIP-based methods, \* marks ViT-based methods and others are CNN-based methods.

## Experiments

### Experimental Setup

**Datasets and Evaluation Protocols.** To fully evaluate the performance of our method, we conduct experiments on three multi-modal object ReID benchmarks. Specifically, RGBNT201 (Zheng et al. 2021) is a multi-modal person ReID dataset comprising RGB, NIR and TIR images. RGBNT100 (Li et al. 2020) is a large-scale multi-modal vehicle ReID dataset with diverse visual challenges, such as abnormal lighting, glaring and occlusion. MSVR310 (Zheng et al. 2022) is a small-scale multi-modal vehicle ReID dataset with more challenges. For evaluation metrics, we utilize the mean Average Precision (mAP) and Cumulative Matching Characteristics (CMC) at Rank-K ( $K = 1, 5, 10$ ). **Implementation Details.** Our model is implemented by using the PyTorch toolbox with one NVIDIA A100 GPU. We employ the pre-trained image encoder of CLIP (Radford et al. 2021) as the backbone. For the input resolution, images are resized to  $256 \times 128$  for RGBNT201 and  $128 \times 256$  for RGBNT100/MSVR310. For data augmentation, we employ random horizontal flipping, cropping and erasing (Zhong et al. 2020). For small-scale datasets (i.e., RGBNT201 and MSVR310), the mini-batch size is set to 64, with 4 images sampled for each identity and 16 identities sampled in a batch. For the large-scale dataset, i.e., RGBNT100, the mini-batch size is set to 128, with 16 images sampled for each identity. We set  $\lambda_1$  to 0.25 and  $\lambda_2$  to 1.0. We use the Adam optimizer to fine-tune the model with a learning rate of  $3.5e^{-4}$ . The warmup strategy with a cosine decay is used for learning rate scheduling. We set the total number of training epochs to 60 for RGBNT201/MVSR310 and 30 for RGBNT100, respectively. The source code is available at <https://github.com/924973292/MambaPro>.

	Methods	RGBNT100		MSVR310	
		mAP	R-1	mAP	R-1
single	PCB (Sun et al. 2018)	57.2	83.5	23.2	42.9
	BoT (Luo et al. 2019)	78.0	95.1	23.5	38.4
	OSNet (Zhou et al. 2019)	75.0	95.6	28.7	44.8
	TransReID* (He et al. 2021)	75.6	92.9	18.4	29.6
	HAMNet (Li et al. 2020)	74.5	93.3	27.1	42.3
	PFNet (Zheng et al. 2021)	68.1	94.1	23.5	37.4
	GAFNet (Guo et al. 2022)	74.4	93.4	-	-
	GPFNet (He et al. 2023)	75.0	94.5	-	-
	CCNet (Zheng et al. 2022)	77.2	96.3	36.4	55.2
	LRMM (Wu et al. 2025)	78.6	<b>96.7</b>	36.7	49.7
Multi	GraFT* (Yin et al. 2023)	76.6	94.3	-	-
	PHT* (Pan et al. 2023)	79.9	92.7	-	-
	UniCat* (Crawford et al. 2023)	79.4	96.2	-	-
	HTT* (Wang et al. 2024b)	75.7	92.6	-	-
	TOP-ReID* (Wang et al. 2023)	81.2	96.4	35.9	44.6
	EDITOR* (Zhang et al. 2024)	82.1	96.4	39.0	49.3
	RSCNet* (Yu et al. 2024b)	82.3	96.6	39.5	49.6
	MambaPro $\ddagger$	<b>83.9</b>	94.7	<b>47.0</b>	<b>56.5</b>

Table 2: Performance on RGBNT100 and MSVR310.

## Performance Comparison

**Multi-modal Person ReID.** As shown in Tab. 1, we compare our proposed MambaPro with state-of-the-art methods on RGBNT201. Within single-modal methods, PCB exhibits a remarkable mAP of 32.8%. For multi-modal methods, CNN-based methods like IEEE (Wang et al. 2022) and LRMM (Wu et al. 2025) achieve competitive performance with mAPs of 49.5% and 52.3%, respectively. However, Transformer-based methods outperform CNN-based methods. Especially, TOP-ReID (Wang et al. 2023) achieves a mAP of 72.3%. Moreover, our MambaPro achieves a 78.9% mAP and 83.4% Rank-1, surpassing TOP-ReID by 6.6% and 6.8%, respectively. These performance improvements verify the effectiveness of our proposed method.

**Multi-modal Vehicle ReID.** As depicted in Tab. 2, single-modal methods generally under-perform multi-modal methods. In single-modal methods, CNN-based methods like BoT (Luo et al. 2019) achieve better results than some Transformer-based methods. For multi-modal methods, LRMM (Wu et al. 2025) achieves a mAP of 78.6% on RGBNT100, showing the effectiveness of its low rank fusion strategy. Meanwhile, Transformer-based methods exhibit their superiority on integrating multi-modal information. Specifically, EDITOR (Zhang et al. 2024) achieves a mAP of 82.1%, surpassing LRMM by 3.5% on RGBNT100. Our MambaPro achieves a mAP of 83.9% on RGBNT100, surpassing EDITOR by 1.8%. Especially on the small-scale dataset MSVR310, our MambaPro achieves a mAP of 47.0%, surpassing EDITOR by 8.0%. These results clearly verify the effectiveness of our proposed methods in multi-modal object ReID tasks.

## Ablation Study

We conduct ablation studies on RGBNT201 to validate the effectiveness of different components. The locked image encoder of CLIP is denoted as CLIP(L) with frozen param-

	Methods					Performance				Complexity	
	CLIP(L)	CLIP(F)	PFA	SRP	MA	mAP	R-1	R-5	R-10	Mem(GB)	Params(M)
A	✓	×	×	×	×	3.3	2.8	7.4	11.1	12.79	0
B	×	✓	×	×	×	69.4	71.5	81.8	87.4	16.85	86.14
C	✓	×	✓	×	×	73.2	76.4	85.2	88.2	15.96	28.34
D	✓	×	×	✓	×	71.3	74.5	83.4	87.9	13.82	28.48
E	✓	×	✓	✓	×	74.4	77.6	87.4	90.3	17.79	56.82
F	✓	×	✓	✓	✓	<b>78.9</b>	<b>83.4</b>	<b>89.8</b>	<b>91.9</b>	20.89	74.20
G	×	✓	✓	✓	✓	78.7	<b>83.5</b>	89.6	<b>92.9</b>	25.22	160.34

Table 3: Performance comparison with key components.

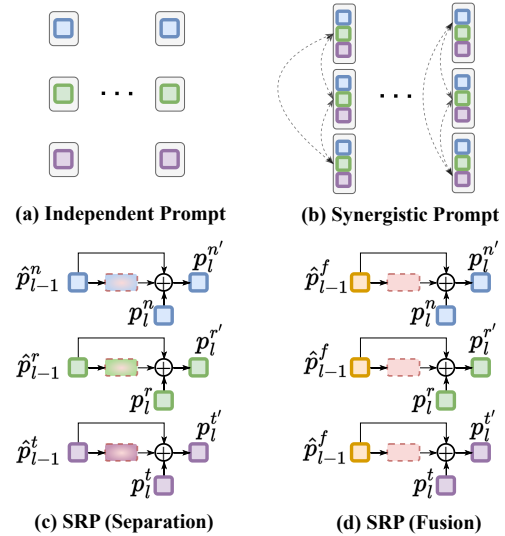


Figure 4: Details of different prompt mechanisms.

eters. Meanwhile, we refer to the fully fine-tuned CLIP as CLIP(F). When MA is not applied, we use  $f_{cls}$  for retrieval. Otherwise, we use  $f_{ma}$  for retrieval.

**Effect of Key Components.** Tab. 3 shows the performance comparison with different components. Comparing Model A with Model B, we observe an impressive mAP improvement by fully fine-tuning CLIP, demonstrating the necessity of fine-tuning. With 32.9% of the trainable parameters compared with Model B, Model C outperforms the fully fine-tuned CLIP by 3.8% mAP, showing the effectiveness of PFA. Meanwhile, Model D also surpasses Model B with minimal additional parameters and FLOPs. Furthermore, through the combination of SRP and PFA, Model E achieves higher performance. Besides, with the efficient modeling of multi-modal long sequences, Model F achieves the best mAP of 78.9%. Model G, with the full fine-tuning of CLIP, achieves a comparable mAP of 78.7%. Comparing the last two rows, our method achieves robust performance no matter whether the backbone is fully fine-tuned or not. With our efficient learning framework, MambaPro transfers the knowledge from CLIP to multi-modal object ReID with superior performance and low complexity.

**Effect of Different Prompt Mechanisms.** Fig. 4 illustrates the details of different prompt mechanisms. Tab. 4 shows the corresponding performances. The first row represents that each modality has its own prompts without synergistic

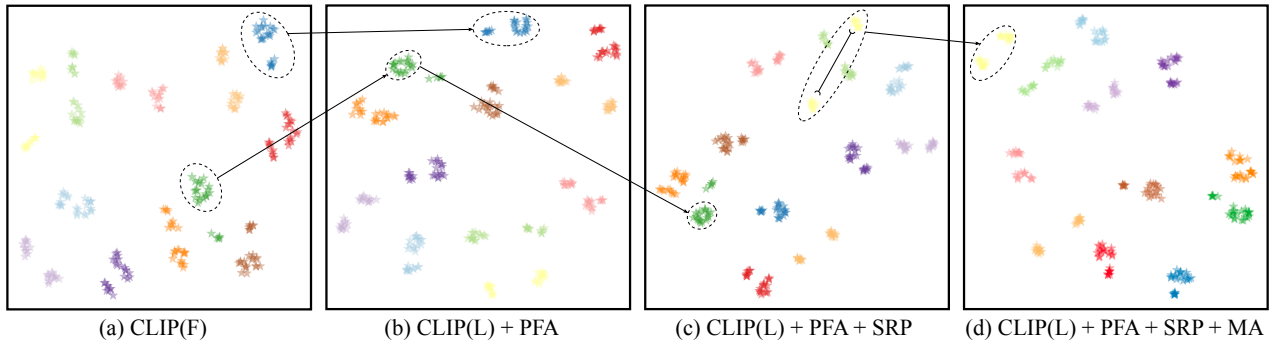


Figure 5: Feature distributions with t-SNE (Van der Maaten and Hinton 2008). Different colors represent different IDs.

Methods	Performance				Prompt Complexity	
	mAP	R-1	R-5	R-10	FLOPs(G)	Params(M)
CLIP(L) + IP	64.7	65.9	77.3	82.9	0.09	0.01
CLIP(L) + SP	68.9	70.7	80.9	85.2	0.28	1.78
CLIP(L) + SRP(Separation)	70.1	72.3	82.1	86.7	0.29	3.55
<b>CLIP(L) + SRP(Fusion)</b>	<b>71.3</b>	<b>74.5</b>	<b>83.4</b>	<b>87.9</b>	0.29	2.37

Table 4: Comparison with different prompt mechanisms.

transfer. SP achieves better performance than Independent Prompt (IP), demonstrating the effectiveness of synergistic transformation between modalities. Both SRP (Separation) and SRP (Fusion) outperform SP, underscoring the effectiveness of residual connections in multi-modal prompts. The difference between SRP (Separation) and SRP (Fusion) lies in the  $\hat{p}_l^f$  comes from the prompts of separate modalities or the fusion of them. SRP (Fusion) performs best by averaging tokens with the residual connections.

**Effect of Different Adapters.** Tab. 5 shows the performance with different adapters. Due to the weaker learning capability without non-linear transformations, LoRA generally achieves inferior performance compared with the Bottleneck Adapter (BNA). Meanwhile, with acceptable complexity, our PFA outperforms BNA by 4.7% mAP, verifying the effectiveness in multi-modal object ReID.

**Effect of Different Aggregation Methods.** In Tab. 6, we compare the performance with different aggregation methods. The Sum and Concat are simple aggregation methods that yield inferior results. The Transformer achieves better performance, demonstrating the effectiveness of attention-based aggregation. Our Intra-MA and Inter-MA outperform Transformers, showcasing the effectiveness of modeling multi-modal sequences with Mamba. Finally, our MA achieves the best performance, demonstrating the effectiveness of our proposed method.

## Visualization

**Feature Distributions.** In Fig. 5, we show the feature distributions with different modules. Comparing Fig. 5 (a) and Fig. 5 (b), CLIP(L) with PFA leads to more compact distributions across different IDs. Besides, with SRP, the gaps between different IDs in Fig. 5 (c) are further expanded. Finally, with MA, the feature distributions become

Methods	Performance				Adapter Complexity	
	mAP	R-1	R-5	R-10	FLOPs(G)	Params(M)
CLIP(L) + LoRA-r(64)	54.5	55.3	68.1	76.1	0.01	0.10
CLIP(L) + LoRA-r(256)	55.7	56.9	69.7	77.3	0.05	0.39
CLIP(L) + LoRA-r(384)	47.2	47.1	62.2	69.1	0.08	0.60
CLIP(L) + BNA	68.5	69.4	79.1	84.0	0.08	0.60
<b>CLIP(L) + PFA</b>	<b>73.2</b>	<b>76.4</b>	<b>85.2</b>	<b>88.2</b>	0.30	2.36

Table 5: Comparison with different adapters. BNA denotes the bottleneck adapter. LoRA-r(x) denotes the LoRA with rank x and we set x to 64, 256 and 384 for comparison.

Methods	Performance				Aggregation Complexity	
	mAP	R-1	R-5	R-10	FLOPs(G)	Params(M)
Sum	75.3	79.4	87.2	89.4	0.01	1.58
Concat	74.3	75.6	84.9	88.9	0.10	2.37
Transformer	76.5	80.2	87.7	90.6	1.36	4.73
Intra-MA	76.2	80.9	88.2	90.5	0.71	5.64
Inter-MA	77.5	81.2	88.6	90.9	0.71	5.41
<b>MA</b>	<b>78.9</b>	<b>83.4</b>	<b>89.8</b>	<b>91.9</b>	1.42	9.47

Table 6: Performance with different aggregation methods.

more separable as shown in Fig. 5 (d). These visualizations strongly support the effectiveness of our proposed modules.

## Conclusion

In this paper, we propose a novel feature learning framework named MambaPro, for multi-modal object ReID. We first employ a Parallel Feed-Forward Adapter (PFA) for knowledge transfer from the pre-trained CLIP to the ReID task. To guide joint learning of multi-modal features, we introduce the Synergistic Residual Prompt (SRP) for interactions of modality-specific prompts. Meanwhile, we propose a Mamba Aggregation (MA) to efficiently aggregate tokenized sequences from different modalities. With linear complexity, the MA efficiently handles long sequences, outperforming existing attention-based methods with low consumption. Extensive experiments on three multi-modal object ReID benchmarks validate the effectiveness of our method and show the potentials of fundamental models.

## Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (No.62101092, 62476044, 62388101), Open Project of Anhui Provincial Key Laboratory of Multimodal Cognitive Computation, Anhui University (No.MMC202102, MMC202407) and Fundamental Research Funds for the Central Universities (No.DUT23BK050, DUT23YG232).

## References

- Ba, J. L.; Kiros, J. R.; and Hinton, G. E. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Chollet, F. 2017. Xception: Deep learning with depthwise separable convolutions. In *CVPR*, 1251–1258.
- Crawford, J.; Yin, H.; McDermott, L.; and Cummings, D. 2023. UniCat: Crafting a Stronger Fusion Baseline for Multimodal Re-Identification. *arXiv preprint arXiv:2310.18812*.
- Diao, H.; Wan, B.; Jia, X.; Zhuge, Y.; Zhang, Y.; Lu, H.; and Chen, L. 2024a. Sherl: Synthesizing high accuracy and efficient memory for resource-limited transfer learning. In *ECCV*, 75–95.
- Diao, H.; Wan, B.; Zhang, Y.; Jia, X.; Lu, H.; and Chen, L. 2024b. Unipt: Universal parallel tuning for transfer learning with efficient parameter and memory. In *CVPR*, 28729–28740.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Elfwing, S.; Uchibe, E.; and Doya, K. 2018. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural networks*, 107: 3–11.
- French, R. M. 1999. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4): 128–135.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- Gu, A.; Goel, K.; and Ré, C. 2021. Efficiently modeling long sequences with structured state spaces. *arXiv preprint arXiv:2111.00396*.
- Gu, A.; Johnson, I.; Goel, K.; Saab, K.; Dao, T.; Rudra, A.; and Ré, C. 2021. Combining recurrent, convolutional, and continuous-time models with linear state space layers. *NeurIPS*, 34: 572–585.
- Guo, J.; Zhang, X.; Liu, Z.; and Wang, Y. 2022. Generative and attentive fusion for multi-spectral vehicle re-identification. In *ICSP*, 1565–1572.
- He, Q.; Lu, Z.; Wang, Z.; and Hu, H. 2023. Graph-Based Progressive Fusion Network for Multi-Modality Vehicle Re-Identification. *TITS*, 1–17.
- He, S.; Luo, H.; Wang, P.; Wang, F.; Li, H.; and Jiang, W. 2021. Transreid: Transformer-based object re-identification. In *ICCV*, 15013–15022.
- Hendrycks, D.; and Gimpel, K. 2016. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*.
- Hermans, A.; Beyer, L.; and Leibe, B. 2017. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*.
- Houlsby, N.; Giurgiu, A.; Jastrzebski, S.; Morrone, B.; De Laroussilhe, Q.; Gesmundo, A.; Attariyan, M.; and Gelly, S. 2019. Parameter-efficient transfer learning for NLP. In *ICML*, 2790–2799.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Ioffe, S.; and Szegedy, C. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 448–456.
- Islam, M. M.; and Bertasius, G. 2022. Long movie clip classification with state-space video models. In *ECCV*, 87–104.
- Islam, M. M.; Hasan, M.; Athrey, K. S.; Braskich, T.; and Bertasius, G. 2023. Efficient Movie Scene Detection using State-Space Transformers. In *CVPR*, 18749–18758.
- Khattak, M. U.; Rasheed, H.; Maaz, M.; Khan, S.; and Khan, F. S. 2023. Maple: Multi-modal prompt learning. In *CVPR*, 19113–19122.
- Lester, B.; Al-Rfou, R.; and Constant, N. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*.
- Li, H.; Li, C.; Zhu, X.; Zheng, A.; and Luo, B. 2020. Multi-spectral vehicle re-identification: A challenge. In *AAAI*, volume 34, 11345–11353.
- Li, S.; Sun, L.; and Li, Q. 2023. CLIP-ReID: exploiting vision-language model for image re-identification without concrete text labels. In *AAAI*, volume 37, 1405–1413.
- Li, Y.; Liu, Z.; Yang, W.; Wang, Y.; Liao, Q.; et al. 2023. CLIP-based Synergistic Knowledge Transfer for Text-based Person Retrieval. *arXiv preprint arXiv:2309.09496*.
- Liang, V. W.; Zhang, Y.; Kwon, Y.; Yeung, S.; and Zou, J. Y. 2022. Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning. *NeurIPS*, 35: 17612–17625.
- 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. *TNNLS*.
- Liu, X.; Zhang, P.; Yu, C.; Lu, H.; and Yang, X. 2021. Watching you: Global-guided reciprocal learning for video-based person re-identification. In *CVPR*, 13334–13343.
- Liu, X.; Zhang, P.; Yu, C.; Qian, X.; Yang, X.; and Lu, H. 2024a. A video is worth three views: Trigeminal transformers for video-based person re-identification. *TITS*, 25: 12818–12828.
- Liu, Y.; Tian, Y.; Zhao, Y.; Yu, H.; Xie, L.; Wang, Y.; Ye, Q.; and Liu, Y. 2024b. Vmamba: Visual state space model. *arXiv preprint arXiv:2401.10166*.
- Luo, H.; Gu, Y.; Liao, X.; Lai, S.; and Jiang, W. 2019. Bag of tricks and a strong baseline for deep person re-identification. In *CVPRW*, 1487–1495.

- Pan, W.; Huang, L.; Liang, J.; Hong, L.; and Zhu, J. 2023. Progressively Hybrid Transformer for Multi-Modal Vehicle Re-Identification. *Sensors*, 23(9): 4206.
- Qian, X.; Fu, Y.; Jiang, Y.-G.; Xiang, T.; and Xue, X. 2017. Multi-scale deep learning architectures for person re-identification. In *ICCV*, 5399–5408.
- 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 *ICML*, 8748–8763.
- Rao, Y.; Chen, G.; Lu, J.; and Zhou, J. 2021. Counterfactual attention learning for fine-grained visual categorization and re-identification. In *ICCV*, 1025–1034.
- Shi, J.; Yin, X.; Chen, Y.; Zhang, Y.; Zhang, Z.; Xie, Y.; and Qu, Y. 2024a. Multi-Memory Matching for Unsupervised Visible-Infrared Person Re-Identification. *arXiv preprint arXiv:2401.06825*.
- Shi, J.; Yin, X.; Zhang, Y.; Xie, Y.; Qu, Y.; et al. 2024b. Learning commonality, divergence and variety for unsupervised visible-infrared person re-identification. In *NeurIPS*.
- Sun, Y.; Zheng, L.; Yang, Y.; Tian, Q.; and Wang, S. 2018. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In *ECCV*, 480–496.
- Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2016. Rethinking the inception architecture for computer vision. In *CVPR*, 2818–2826.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *JMLR*, 9(11).
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *NeurIPS*, 30.
- Wan, Z.; Zhang, P.; Wang, Y.; Yong, S.; Stepputtis, S.; Sycara, K.; and Xie, Y. 2024. Sigma: Siamese mamba network for multi-modal semantic segmentation. *arXiv preprint arXiv:2404.04256*.
- Wang, Y.; Liu, X.; Zhang, P.; Lu, H.; Tu, Z.; and Lu, H. 2023. TOP-ReID: Multi-spectral Object Re-Identification with Token Permutation. *arXiv preprint arXiv:2312.09612*.
- Wang, Y.; Zhang, P.; Wang, D.; and Lu, H. 2024a. Other tokens matter: Exploring global and local features of Vision Transformers for Object Re-Identification. *CVIU*, 244: 104030.
- Wang, Z.; Huang, H.; Zheng, A.; and He, R. 2024b. Heterogeneous Test-Time Training for Multi-Modal Person Re-identification. In *AAAI*, volume 38, 5850–5858.
- Wang, Z.; Li, C.; Zheng, A.; He, R.; and Tang, J. 2022. Interact, embed, and enlarge: Boosting modality-specific representations for multi-modal person re-identification. In *AAAI*, volume 36, 2633–2641.
- Wu, D.; Liu, Z.; Chen, Z.; Gan, S.; Tan, K.; Wan, Q.; and Wang, Y. 2025. LRMM: Low rank multi-scale multi-modal fusion for person re-identification based on RGB-NI-TI. *ESWA*, 263: 125716.
- Xu, Z.; Tang, F.; Chen, Z.; Zhou, Z.; Wu, W.; Yang, Y.; Liang, Y.; Jiang, J.; Cai, X.; and Su, J. 2024. Polyp-Mamba: Polyp Segmentation with Visual Mamba. In *MICCAI*.
- Yang, B.; Chen, J.; and Ye, M. 2023. Towards Grand Unified Representation Learning for Unsupervised Visible-Infrared Person Re-Identification. In *ICCV*, 11069–11079.
- Yang, B.; Chen, J.; and Ye, M. 2024. Shallow-Deep Collaborative Learning for Unsupervised Visible-Infrared Person Re-Identification. In *CVPR*, 16870–16879.
- Yang, L.; Kang, B.; Huang, Z.; Xu, X.; Feng, J.; and Zhao, H. 2024. Depth anything: Unleashing the power of large-scale unlabeled data. *arXiv preprint arXiv:2401.10891*.
- Yin, H.; Li, J.; Schiller, E.; McDermott, L.; and Cummings, D. 2023. GRAFT: Gradual Fusion Transformer for Multi-modal Re-Identification. *arXiv preprint arXiv:2310.16856*.
- Yu, C.; Liu, X.; Wang, Y.; Zhang, P.; and Lu, H. 2024a. TF-CLIP: Learning text-free CLIP for video-based person re-identification. In *AAAI*, volume 38, 6764–6772.
- Yu, Z.; Huang, Z.; Hou, M.; Pei, J.; Yan, Y.; Liu, Y.; and Sun, D. 2024b. Representation Selective Coupling via Token Sparsification for Multi-Spectral Object Re-Identification. *TCSVT*.
- Zhang, G.; Zhang, P.; Qi, J.; and Lu, H. 2021. Hat: Hierarchical aggregation transformers for person re-identification. In *ACM MM*, 516–525.
- Zhang, P.; Wang, Y.; Liu, Y.; Tu, Z.; and Lu, H. 2024. Magic tokens: Select diverse tokens for multi-modal object re-identification. In *CVPR*, 17117–17126.
- Zheng, A.; He, Z.; Wang, Z.; Li, C.; and Tang, J. 2023. Dynamic Enhancement Network for Partial Multi-modality Person Re-identification. *arXiv preprint arXiv:2305.15762*.
- Zheng, A.; Wang, Z.; Chen, Z.; Li, C.; and Tang, J. 2021. Robust multi-modality person re-identification. In *AAAI*, volume 35, 3529–3537.
- Zheng, A.; Zhu, X.; Ma, Z.; Li, C.; Tang, J.; and Ma, J. 2022. Multi-spectral vehicle re-identification with cross-directional consistency network and a high-quality benchmark. *arXiv preprint arXiv:2208.00632*.
- Zhong, Z.; Zheng, L.; Kang, G.; Li, S.; and Yang, Y. 2020. Random erasing data augmentation. In *AAAI*, volume 34, 13001–13008.
- Zhou, K.; Yang, Y.; Cavallaro, A.; and Xiang, T. 2019. Omni-scale feature learning for person re-identification. In *ICCV*, 3702–3712.
- Zhu, L.; Liao, B.; Zhang, Q.; Wang, X.; Liu, W.; and Wang, X. 2024. Vision mamba: Efficient visual representation learning with bidirectional state space model. *arXiv preprint arXiv:2401.09417*.