

# All in One Framework for Multimodal Re-identification in the Wild

He Li<sup>1</sup> Mang Ye<sup>1\*</sup> Ming Zhang<sup>2</sup> Bo Du<sup>1</sup>

<sup>1</sup>National Engineering Research Center for Multimedia Software, Institute of Artificial Intelligence, School of Computer Science, Hubei LuoJia Laboratory, Wuhan University, Wuhan, China,

<sup>2</sup>Guangzhou Urban Planning Design Survey Research Institute, Guangzhou, China

## Abstract

In Re-identification (ReID), recent advancements yield noteworthy progress in both unimodal and cross-modal retrieval tasks. However, the challenge persists in developing a unified framework that could effectively handle varying multimodal data, including RGB, infrared, sketches, and textual information. Additionally, the emergence of large-scale models shows promising performance in various vision tasks but the foundation model in ReID is still blank. In response to these challenges, a novel multimodal learning paradigm for ReID is introduced, referred to as All-in-One (AIO), which harnesses a frozen pre-trained big model as an encoder, enabling effective multimodal retrieval without additional fine-tuning. The diverse multimodal data in AIO are seamlessly tokenized into a unified space, allowing the modality-shared frozen encoder to extract identity-consistent features comprehensively across all modalities. Furthermore, a meticulously crafted ensemble of cross-modality heads is designed to guide the learning trajectory. AIO is the **first** framework to perform all-in-one ReID, encompassing four commonly used modalities. Experiments on cross-modal and multimodal ReID reveal that AIO not only adeptly handles various modal data but also excels in challenging contexts, showcasing exceptional performance in zero-shot and domain generalization scenarios. Code will be available at: <https://github.com/lihe404/AIO>.

## 1. Introduction

Person Re-identification (ReID) aims to retrieve a target person captured by multiple non-overlapping cameras [69]. It is widely used in intelligent surveillance, security, and many other fields. ReID has been deeply studied in recent years and achieves human-level performance in both unimodal and cross-modal retrieval tasks [1, 21, 68, 77].

Existing works are capable of retrieving between RGB

\*Corresponding Author: Mang Ye (yemang@whu.edu.cn)

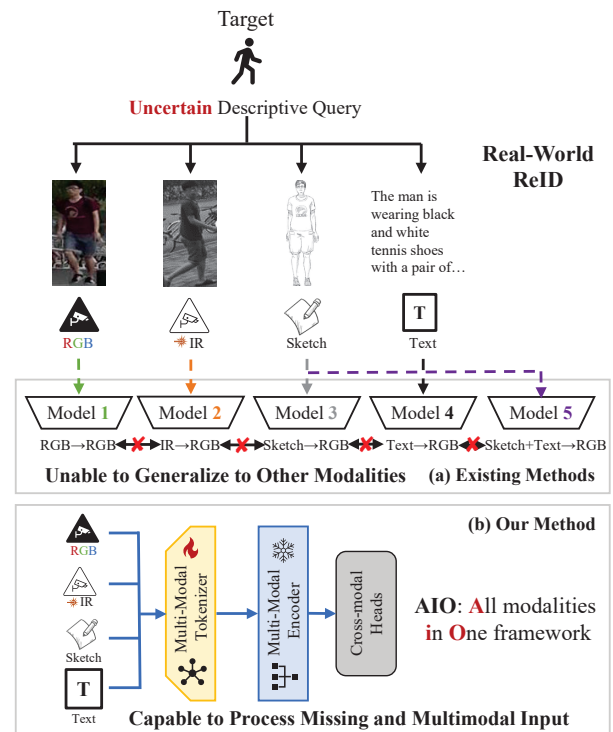


Figure 1. **Illustration of the proposed AIO and existing methods.** (a) Existing ReID methods [7, 68, 86] independently learn the cross-modal ReID models, incapable of handling the uncertain input modalities in real-world scenarios. (b) Our proposed AIO framework exhibits the capability to proficiently manage diverse combinations of input modalities, thus addressing the inherent uncertainties prevalent in practical deployment scenarios.

images or leveraging different modalities of the query (e.g. infrared (IR), sketch, or text) to find the person in RGB images [13, 22, 23, 64, 72]. However, RGB images are susceptible to environmental light fluctuations, while IR and sketch images lack vital color information crucial for ReID tasks. The adage “a picture is worth a thousand words” underscores the ease of accessing textual information, though it falls short in providing intricate visual details [51]. Moreover, as shown in Fig. 1(a) and Tab. 1, existing cross-modal

Method	RGB	IR	Sketch	Text	Multi
NFormer [55]	✓	✗	✗	✗	✗
DC-Former [32]	✓	✗	✗	✗	✗
AGW [69]	✓	✓	✗	✗	✗
DART [65]	✓	✓	✗	✗	✗
Gui <i>et al.</i> [17]	✓	✗	✓	✗	✗
SketchTrans [6]	✓	✗	✓	✗	✗
APTM [66]	✓	✗	✗	✓	✗
BiLMa [13]	✓	✗	✗	✓	✗
TriReID [75]	✓	✗	✓	✓	✓(3)
UNIReID [7]	✓	✗	✓	✓	✓(3)
AIO	✓	✓	✓	✓	✓(4)

Table 1. **Comparison of AIO and existing methods on cross-/multi-modality retrieval.** The number in (·) in the “Multi” column indicates the number of support modalities at inference time.

methods are confined to specific paired modalities and models during training, rendering them unable to handle diverse input modalities effectively. Consequently, the generalizability of these methods to unseen modalities is severely hampered, a significant hurdle given the uncertainty of the modality in real-world query scenarios. This lack of adaptability across modalities severely constrains the practical applicability of existing methods in a practical real-world deployment with uncertain testing environments. Thus, 1) *How to improve the generalizability of modality is a significant challenge?*

Meanwhile, in real-world scenarios, individuals of interest frequently encounter unknown environments that are not learned during training, *i.e.*, zero-shot ReID in the wild. Existing methods explore domain generalizability based on a single modality, which fails to handle multi-modal zero-shot retrieval. Recently, foundation large models have shown their power in diverse language and vision tasks. Pioneering models such as CLIP [48] and CoCa [73], exemplify the prowess of large-scale pre-trained foundation models as robust zero-shot performers. This characteristic holds significant relevance for ReID tasks. Despite the existence of several large-scale ReID pre-trained models [11, 41, 66], their zero-shot performance falls short of expectations. Typical down-stream fine-tuning or training strategies would be too resource-demanding in a new challenging scenarios, *e.g.*, data collection and annotation. Moreover, the cost of training a large-scale foundation model is too high to afford for most researchers and small companies. Thus, 2) *Is there a straightforward method to utilize extensive pre-trained foundational models for improving zero-shot performance in ReID with uncertain modalities?*

To address the aforementioned issues, we introduce an innovative All-in-One (AIO) framework to tackle the challenges inherent in zero-shot multimodal ReID. As shown in Fig. 1(b), the key idea of our work is to explore the potential of leveraging transformer-based foundation models to

address uncertain multimodal retrieval, enhancing zero-shot ability in multimodal ReID, *e.g.* any combination of RGB, IR, sketch, text, or simple cross-modal retrieval. AIO represents an experimental effort, being the **first** framework capable of simultaneously accommodating all four commonly used modalities in different ReID tasks concurrently.

In order to achieve the above goal, AIO firstly designs a lightweight multimodal tokenizer to unify diverse data. It is followed by a frozen foundation model that serves as a shared feature encoder, extracting a generalized semantic representation across all modalities and improving zero-shot performance. Then, to guide cross-modal and multimodal feature learning, AIO proposes several cross-modal heads which contain: *a) A Conventional Classification Head* is utilized as the foundation of the learning guidance, learning identity-invariant representations; *b) Vision Guided Masked Attribute Modeling* is introduced to learn fine-grained features and build a relationship between text and images; *c) Multimodal Feature Binding* is utilized to close features of diverse modalities together.

Furthermore, the acquisition of multimodal data in real-world scenarios poses considerable challenges, particularly concerning IR and Sketch images. The shortage of IR cameras and the substantial human labor involved in sketch drawing contribute to this difficulty. Existing multimodal learning methods [47, 57] demand paired multimodal, a pre-requisite not consistently met in realistic environments. In addressing the absence of certain modalities in multimodal learning, except the proposed Multimodal Feature Binding, integrates synthetic augmentation strategies, CA [72] and Lineart [53], to generate synthetic IR and Sketch images respectively. CA and Lineart have been shown to effectively diminish the domain gap between RGB-IR and RGB-Sketch modalities. Their utility extends to acting as a bridge that connects feature representations of the same target across diverse modalities, thereby facilitating the reduction of the modality gap [54, 72].

Comprehensive experiments across various zero-shot cross-modal and multimodal ReID scenarios, involving all four modalities, are conducted to evaluate the performance of the proposed framework. We also explore different foundation models and multimodal input data combinations to assess the versatility of AIO. The proposed framework demonstrates remarkable performance on multimodal ReID tasks and competitive performance on cross-modal ReID tasks without additional fine-tuning, highlighting its potential as a robust zero-shot solution for complex multimodal ReID tasks. In summary, our contributions are three-fold:

- We identify a critical limitation in existing cross-modal ReID that they lack generalizability to novel modalities, coupled with poor zero-shot performance.
- We design a novel multimodal ReID framework, which innovatively integrates a pre-trained foundation model

and a multimodal tokenizer into ReID tasks, complemented a missing modality synthesis strategy, and three cross-modal heads to learn a unified multimodal model.

- We perform extensive and comprehensive experiments to demonstrate that our AIO framework is capable of handling uncertain input modalities for ReID tasks, achieving competitive performance on zero-shot cross-modal and multi-modal ReID tasks.

## 2. Related Work

### 2.1. Cross-modal ReID

Person Re-identification (ReID) can be classified into single-modal [18, 27, 40, 79] ReID and cross-modal ReID [7, 62]. Specifically, cross-modal ReID considers special cases in which RGB images of the target are not available but non-RGB modalities, such as infrared (IR) [63, 65, 68, 71], sketch [6, 17], or description of the person [15, 31, 51, 66, 86] could be leveraged to enlarge the application area of ReID technology.

In real-world scenarios, the RGB image of the target may not be directly available. On the one hand, the crime often happens at night, when RGB cameras cannot capture high-quality images for ReID but IR cameras could give relatively good images of the person [69]. Thus, adopting IR images for target retrieval enters the picture. Zhang *et al.* [77] proposes a feature-level modality compensation network to compensate for the missing modality-specific information at the feature level to help the model learn discriminative features. Wu *et al.* [61] leverage unlabeled data and propose a progressive graph matching method to learn a relationship between IR and RGB images to alleviate the high cost of annotation. On the other hand, the natural language descriptions of the witness or the sketch drawn based on the textual descriptions are easy to access. Pang *et al.* [46] first design an adversarial learning method to learn domain-invariant features cross sketch and RGB. Zhai *et al.* [75] introduce a multimodal method that combines both sketch and text as queries for retrieval.

However, in the real world, the modality of the given target information is uncertain. The aforementioned works could only handle exactly two modalities, which hinders the applicability of these methods. To alleviate this problem, Chen *et al.* [7] proposes a modality-agnostic retrieval method that leverages RGB, sketch, and text modalities to learn modality-specific features and fuse them based on different uni/multimodal tasks. The proposed method could handle any combination of three learned modalities, expanding usage scenarios and reducing limitations. Nevertheless, it does not take IR into account and requires a complex design that lacks scalability. Different from that, the architecture of the proposed AIO is simple and it is easy to be extended to more modalities.

### 2.2. Multimodal Learning

Multimodal learning methods aim to utilize the complementary properties of various modalities to learn the semantics of a task [34]. Recently, multimodal transformers [14, 24, 29] have emerged as unified models that fuse different modality inputs with token concatenation rather than extracting modality-specific and cross-modality representations. However, most multimodal learning methods [3, 47, 57] are designed based on the assumption of the completeness of modality for training or inference, which is not always held in the real world. To face this challenge, some researchers [4, 25, 45, 56] explore building multimodal methods that could handle missing modalities. ImageBind [16] projects all features of different modalities into the same feature space and leverages contrastive learning to align all modalities to a based modality. SMIL [42] proposes to estimate the latent features of the missing modality via Bayesian Meta-Learning. GCNet [33] introduces graph neural network-based modules to capture temporal and speaker dependencies and jointly optimize classification and reconstruction tasks. Similar to previous work, the proposed AIO aspires to project features from diverse modalities into a unified feature embedding. Notably, we capitalize on the inherent capabilities of the transformer, adept at handling variable input lengths. This strategy enables the model to seamlessly accommodate inputs originating from any combination of modalities, enhancing the flexibility and adaptability of the AIO framework.

### 2.3. Foundation Model

Foundation models are designed to be adapted to various downstream tasks by pre-training on broad data at scale [2]. The efficacy of large-scale pre-trained models is evident in their capacity to enhance data encoding and elevate the performance of downstream tasks [20, 74]. Recent investigations [43, 49, 50, 67] reveal the notable emergent ability exhibited by most large-scale pre-trained foundation models, particularly in the context of robust zero-shot performance. CLIP [48] focuses on multimodal contrastive learning on noisy web image-text pairs to learn aligned image and text representation. Impressively, CLIP achieves accuracy comparable to the original ResNet-50 on ImageNet zero-shot, without exposure to any samples from ImageNet. DALL-E [49] introduces a simple approach that autoregressively models sufficient and large-scale text and image tokens and demonstrates commendable zero-shot performance when compared to preceding domain-specific models. Nonetheless, the foundation model within ReID remains in a nascent state. The suboptimal zero-shot performance of extant large-scale pre-trained models persists due to challenges in data acquisition. Motivated by the impressive zero-shot capabilities exhibited by foundation models, the proposed AIO framework strategically employs frozen

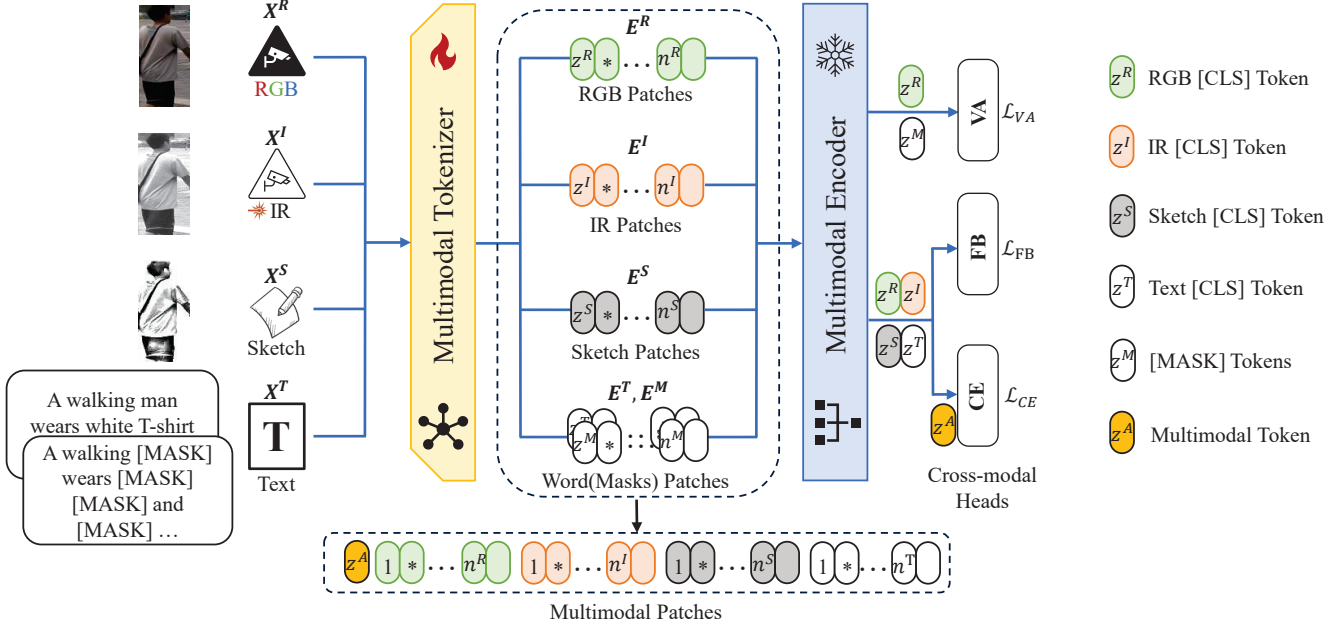


Figure 2. **The schematic of the proposed AIO framework.** VA: Vision Guided Masked Attribute Modeling head, FB: Feature Binding head, CE: Classification head. Our framework mainly contains three parts: I) a learnable multimodal tokenizer to project diverse modalities into a unified embedding, II) a frozen foundation modal to extract complementary cross-modal representations, and III) several cross-modal heads used to dig cross-modality relationships. In order to alleviate the missing modality problem, we also leverage Channel Augmentation [72] and Linear [53] to synthesize IR and sketch images that are missing.

pre-trained large-scale foundation models as feature extractors, which aims to imbue our framework with the ability to learn generalized semantics from a diverse range of modalities, thereby enhancing its zero-shot performance.

### 3. All in One Framework

In this section, we describe the proposed AIO framework in detail. AIO exhibits the capability to adeptly handle uncertain multimodal data, encompassing RGB, IR, Sketch, and Text modalities. To realize this, we introduce a novel multimodal tokenizer designed to project data into a unified embedding space. Simultaneously, a large-scale pre-trained foundation model serves as the shared feature extractor, encoding embeddings from diverse modalities. The learning process of the multimodal tokenizer is guided by cross-modal heads that are designed specifically for ReID tasks. The schematic of the AIO is illustrated in Fig. 2.

**Preliminary.** Formally, within the AIO framework, three key components are denoted as follows: I) the multimodal tokenizer  $\psi^{mod}(\cdot)$ , II) the frozen multimodal encoder  $f(\cdot)$ , and III) cross-modal heads  $\Upsilon_{head}(\cdot)$ , where  $mod$  refers to the notation of each modality, such as RGB (R), IR (I), Sketch (S), and Text (T);  $head$  refers to the notation of Classification (CE), Vision Guided Masked Attribute Modeling (VA), and Feature Binding head (FB), respectively. The inputs are denoted as  $x^{mod} \in X^{mod}$ , the embeddings generated by tokenizers are  $E^{mod}$ , and the output feature from the frozen multimodal encoder is  $z^{mod}$ ,

which is corresponding to class tokens in Fig. 2. We assume 1) each modality possesses a specific parameter space  $\theta^{mod}$  for modality-specific feature representations, and 2) there exists a shared parameter space  $\theta^A$ , an intersection of each modality parameter space for modality-shared feature representations, adhering to the condition:

$$\theta^A \in \theta^R \cap \theta^I \cap \theta^S \cap \theta^T \text{ and } \theta^A \neq \emptyset; \quad (1)$$

For a simple multimodal network with a cross-modal head  $\Upsilon$  can be written as:

$$\theta^A = \arg \min_{x \in X^{mod}} \{\Upsilon_{head} \circ z^{mod}\} \quad (2)$$

where  $z = f \circ \psi^{mod}(x)$ ,  $x \in X^{mod}$  is the input from any modality,  $\circ$  is the function composition operation. The notation of different modalities used in AIO is presented in Tab. 2.

#### 3.1. Multimodal Tokenizer

To project various modalities into a unified space, we devise a straightforward multimodal tokenizer. This tokenizer comprises four projectors: three dedicated to RGB, IR, and Sketch modalities, and one for Text. Furthermore, a multimodal embedding is constructed by amalgamating the embeddings from the respective modalities.

**Image Tokenizers.** Given the disparate channel counts in RGB, IR, and Sketch images, for the sake of convenience, we employ channel replication in both IR and

Source	Modality	Input	Embedding	Feature
Raw	RGB (R)	$X^R$	$E^R$	$Z^R$
	IR (I)	$X^I$	$E^I$	$Z^I$
	Sketch (S)	$X^S$	$E^S$	$Z^S$
	Text (T)	$X^T$	$E^T$	$Z^T$
	Multimodal (A)	-	$E^A$	$Z^A$
Synthesis	CA [72]	$X^I$	$E^I$	$Z^I$
	Lineart [53]	$X^S$	$E^S$	$Z^S$
	Masked Text	$X^M$	$E^M$	$Z^M$

Table 2. **Notation of different modalities.** Be aware that both Lineart [53] and CA [72] operations are introduced, and the generated images are considered as Sketch and IR images when these two modalities are missing.

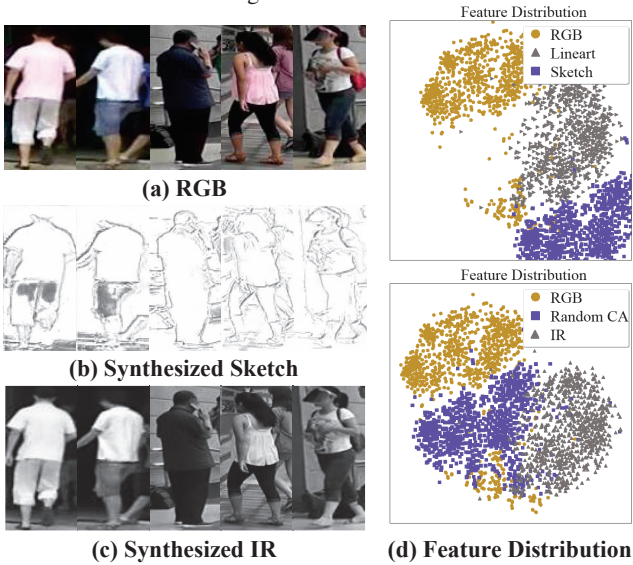


Figure 3. **The generated synthetic Sketch and IR images.** We also visualize the feature distribution of RGB, IR, Sketch, and synthesized images.

Sketch modalities to align their channel count with the three channels present in RGB images. Deviating from the original tokenizer employed in ViT [10], leading to induce training instability [60], we opt for the IBN [44] style tokenizer from ViT-ICS [41]. The convolutional, batch normalization (BN), and rectified linear unit (ReLU) layers inherent to the IBN-style tokenizer substantially enhance training stability [5] and mitigate data bias which is critical for ReID [41].

**Text Tokenizers.** In accordance with prior research efforts [23], we adopt the CLIP tokenizer [48] to directly map the text. Each word is uniquely associated with a token, and through the utilization of word embedding layers, it is projected into a high-dimensional feature space to yield a sequence of word embeddings.

**Multimodal Embedding.** In the context of multimodal embedding, the embeddings originating from various modalities are concatenated. Additionally, following previous works [8, 10], a learnable token  $z^A$  is appended to the sequence of multimodal embeddings. Simultaneously, posi-

tion embeddings  $E^{Pos}$  are employed to enhance position information, seamlessly integrated with the multimodal embeddings via element-wise addition, a procedure akin to the original operation in ViT [10]. The multimodal embedding is formulated as follows:

$$E^A = [z^A, E^R, E^I, E^S, E^T] + E^{Pos}, \quad (3)$$

$$E \in \mathbf{R}^{n \times D}, E^{Pos} \in \mathbf{R}^{(n+1) \times D}.$$

### 3.2. Missing Modality Synthesis

Given the insufficiency of multimodal data in ReID, especially in IR, and Sketch, we introduce Channel Augmentation (CA) [72] and Lineart [53] as augmentation methods to synthesize absent modalities. The generated samples are shown in Fig. 3(a)-(c). The incorporation of synthetic modalities offers two advantages: 1) an expansion in the size of the multimodal sample, thereby mitigating issues associated with missing modalities; 2) CA and Lineart act as conduits bridging the gap between synthetic and real IR and sketch modalities. This is attributed to the feature distribution of the augmented images aligning between RGB and real IR and Sketch images. The visual representation of the feature distribution for RGB-Lineart-Sketch and RGB-CA-IR, as presented in Fig. 3(d) through t-SNE, serves as evidence of their efficacy in alleviating the learning challenges arising from modality gaps.

**Progressively Learning with Synthetic data.** We employ a progressively learning strategy to train the proposed AIO framework. The strategy involves initially training on synthetic images, incorporating real-world RGB and Text, for a few number of epochs. Subsequently, the model undergoes further fine-tuning using paired IR and Sketch images from the real world. This sequencing is deliberate, as synthetic images exhibit a reduced domain gap with RGB compared to real IR and Sketch images, facilitating a more accessible learning process for the model. A similar phenomenon is also found in other cross-modal works [19, 39, 70, 76, 83].

### 3.3. Multimodal Modeling and Binding

All representations extracted by the frozen multimodal encoder from each embedding are fed into cross-modal heads  $\Upsilon_{head}$ , which are specifically designed to learn cross-modal relationships between different modalities. As illustrated in Fig. 2, there are three heads: 1) *Conventional Classification Head*, learning identity invariant representations like in other ReID works [26, 35, 87]; 2) *Vision Guided Masked Attribute Modeling*, seeking to learn fine-grained RGB-Text relationships; 3) *Multimodal Feature Binding*, aiming to align each modality representations together.

**Conventional Classification (CE).** The classification head only contains a bottleneck [40] and a classifier, which is

Partition	Dataset	Venue	#ID	#RGB Imgs	#IR Imgs	#Sketch Imgs	#Text
Train	SYNTH-PEDES [88]	arXiv23	312,321	4,791,771	-	-	12,138,157
	LLCM [78]	CVPR23	1,064	25,626	21,141	-	-
	MaSk1K [37]	ACMMM23	996	32,668	-	4,763	-
Test	Market1501 [82]	ICCV15	1,501	32,668	-	-	-
	SYSU-MM01 [59]	IJCV20	491	30,071	15,792	-	-
	PKU-Sketch [46][75]	ACMMM18	200	400	-	200	200
	CUHK-PEDES [30]	CVPR17	13,003	40,206	-	-	80,412
	Tri-CUHK-PEDES [7]	CVPR23	13,003	40,206	-	40,206	80,412

Table 3. **The statistics of datasets used in experiments.** More details can be found in corresponding papers.

constrained by Cross-Entropy loss as follows:

$$\mathcal{L}_{CE} = -\frac{1}{N} \sum y \log(\Upsilon_{CE} \circ z^{mod}), \quad (4)$$

where  $N$  is the number of pedestrian IDs,  $\Upsilon_{CE}$  indicates the Conventional Classification head. The conventional Triplet Loss, commonly employed in related frameworks, is omitted in our architecture, as we opt for the utilization of a multimodal feature binding loss.

**Vision Guided Masked Attribute Modeling (VA).** Attributes play a pivotal role in highlighting essential characteristics of an individual, encompassing factors such as gender and hair color. These attributes are instrumental in cross-modal alignment and the differentiation of distinct individuals. In this context, we investigate the utility of attribute information embedded in the Text modality to serve as supervisory signals for learning discriminative person representations. To be specific, in the case of a paired RGB image and Text, we adopt a strategy from the prior work [66], where specific attribute keywords in the Text are selectively masked. These masked words are then projected to a special token  $[MASK]$ . Subsequently, the concatenated features of the paired RGB image and the masked token are fed into a decoder structured with MLPs and a classifier, represented as follows:

$$\mathcal{L}_{VA} = -\frac{1}{N_A M} \sum \sum y \log(\Upsilon_{VA}(z^R \oplus z^m)), \quad (5)$$

where  $\Upsilon_{VA}$  indicates the Vision Guided Masked Attribute Modeling head,  $z^m \in z^M$  are the features of masked tokens,  $N_A, M$  are the number of classes and the number of masked tokens, and  $\oplus$  denotes the concatenation operation.

**Multimodal Feature Binding (FB).** To align all modalities onto a shared manifold, we attract features from all modalities towards the RGB feature. This alignment is facilitated through the incorporation of a novel supervised feature binding loss, elucidated in the subsequent section:

$$\mathcal{L}_{FB} = -\sum \log \frac{\exp(\frac{1}{mod} \sum_{mod \neq R} \|z_i^R, z_i^{mod}\|/\tau)}{\sum_{i \neq j} \exp(\|z_i^R, z_j^R\|/\tau)}, \quad (6)$$

where  $\|\cdot\|$  is the cosine similarity,  $z_i^R$  is the representation of person  $i^{th}$  RGB embedding,  $z_i^{mod}$  are the representations of person  $i^{th}$  other modalities embeddings,  $z_j^R$  are the RGB representations belonging to other people,  $\tau$  is the temperature that controls the smoothness of the softmax distribution. Diverging from the conventional InfoNCE approach [52], our feature binding loss involves bringing together features from all modalities corresponding to the same individual, while simultaneously creating a separation between RGB features of distinct individuals, rather than applying the same principle to all features. This difference is motivated by the prevalence of RGB as the most common modality in real-world scenarios, contributing the most abundant data and consistently present in all publicly available datasets.

### 3.4. Overall Architecture

As elucidated earlier, the primary objective of the AIO framework is to learn a multimodal tokenizer through a frozen multimodal encoder, under the guidance of cross-modal heads. We believe that the emergent capabilities demonstrated in large-scale foundation models can effectively augment the zero-shot ability in multimodal ReID tasks. Additionally, capitalizing on the inherent adaptability of transformer architecture to accommodate variable input lengths, AIO exhibits competence in processing diverse combinations of commonly employed modalities in ReID. To realize this objective, AIO is constrained by three cross-modal heads, that can be written as follows:

$$\mathcal{L} = \mathcal{L}_{CE} + \alpha \mathcal{L}_{VA} + \mathcal{L}_{FB}, \quad (7)$$

where  $\alpha$  is a fixed weight to control the importance of Vision Guided Masked Attribute Modeling.

## 4. Experiment

In this section, we conduct a comprehensive evaluation of the proposed AIO framework across both cross-modal and multimodal ReID tasks. Our analysis demonstrates the efficacy of the AIO framework, particularly in zero-shot scenarios involving uncertain input modalities within ReID tasks. Additionally, we delve into the examination of varying foundation models and input modality combinations.

Heads	R→R	I→R	S→R	T→R	S+T→R
Base	3.2	1.5	1.6	21.6	0.5
+ VA	6.4	0.9	1.2	52.4	0.6
+ CE	78.0	43.7	46.8	53.5	91.8
+ FB	79.6	57.6	70.2	53.4	92.1

Table 4. **Effectiveness of each cross-modal heads.** The Rank-1 zero-shot performance is reported.

**The short notation of each modality will be used in this section,** details can be found in Tab. 2.

#### 4.1. Experimental settings

**Datasets.** Three publicly available datasets SYNTH-PEDES [88] for R-T pairs, LLCM [78] for R-I images, MaSk1K [37] for R-S images are leveraged for training. For zero-shot performance evaluation, five widely used real-world datasets are used for evaluations, Market1501 [82] for R→R task, SYSU-MM01 [59] for I→R task, PKU-Sketch [46] for S→R task, CUHK-PEDES [30] for T→R task, and Tri-CUHK-PEDES [7] for T+S→R task. The dataset statistics are shown in Tab. 3. More details can be found in the original papers.

**Evaluation Protocols.** Following existing cross-modality ReID settings [6, 31, 36, 61], we use the Rank- $k$  matching accuracy, mean Average Precision (mAP) metrics, and mean Inverse Negative Precision (mINP) [69] for performance assessment. In the context of multimodal ReID, we adhere to the evaluation settings outlined in TriReID [75] and UNIREID [7] specifically designed for RGB-Text+Sketch scenarios. To accommodate other multimodal data combinations, we leverage CA [72] and Lineart [53] to generate simulated IR and Sketch images. While acknowledging that this may not perfectly simulate real-world scenarios, it provides valuable insights into the multimodal performance of the proposed AIO framework.

**Implementation Details.** We employ the ViT [10] as the backbone, which is pre-trained on LAION-2B dataset with contrastive learning, to reinforce the ability for generic token encoding. All parameters of the backbone networks are frozen. The Text tokenizer is from the pre-trained CLIP [48] to segment sentences into subwords and transform them into word embeddings. We perform a progressively learning strategy training process in AIO framework, as we discussed in Sec. 3.2. *stage 1*) In the first 40 epochs, we sample 32 paired RGB and text samples from SYNTH-PEDES only combined with generated synthetic IR and Sketch images using CA [72] and Lineart [53]. Moreover, we randomly chose two to four embeddings from different modalities to build the multimodal embedding. It is worth noting that, multimodal embedding may not contain RGB embedding. *stage 2*) In the rest 80 epochs, we still select 32 samples for a batch but from all training datasets. For data from SYNTH-PEDES, the sampling, synthetic methods, and construction of multimodal embedding are unchanged. For data from

Model	R→R	I→R	S→R	T→R	S+T→R
ViT* [41]	74.2	49.6	63.4	48.2	79.2
Uni* [28]	76.3	51.3	68.1	52.6	84.3
CLIP* [48]	77.7	55.5	69.5	52.8	87.6
LAION [80] †	79.6	57.6	70.2	53.4	92.1

Table 5. **Zero-shot performance of AIO with different foundation models.** The Rank-1 performance is reported. \* indicates the tokenizer of the original model is replaced by ours. † is the backbone used in AIO.

Multimodal	Rank-1	Rank-5	Rank-10
R+T	56.5	76.2	85.1
R+I	48.2	70.7	79.3
I+T	53.4	74.3	81.0
I+S	48.1	70.6	79.0
R+I+T	57.8	78.0	86.3
I+S+T	55.6	74.1	82.2
R+I+S+T	58.6	77.9	86.6

Table 6. **Zero-shot performance with multimodal input on Tri-CUHK-PEDES.** Be aware that the IR images are generated by using CA [72] rather than real-world IR images.

LLCM and MaSk1K, only paired RGB-IR and RGB-Sketch images are leveraged. The multimodal embedding for samples from these two datasets only contains available modalities. We also apply random horizontal flipping and random cropping for visual modalities. All images are resized to  $384 \times 192$ . The framework is optimized by AdamW [38] optimizer with a base learning rate of  $1e-4$ , a cosine weight decay of  $1e-4$ , and a warmup in the first 5 epochs. The learning rate of the CLIP tokenizer is multiplied by  $1e-1$  since they have already been pre-trained. The  $\alpha$  is set to  $3e-1$  and the  $\tau$  is set to  $5e-2$  as in [16]. The framework is distributively trained on 8 NVIDIA 3090 GPUs.

#### 4.2. Ablation Study

**Efficacy of Designed Modules.** We first evaluate the effectiveness of the designed components. As evident from Tab. 4, each introduced cross-modal head proves crucial for the overall performance of our All-in-One (AIO) framework. Specifically, VA head yields the most substantial performance enhancement in the T→R task, CE head plays an important role in the R→R task, and FB head improves all cross-modal and multimodal tasks.

**Different Foundation Models.** We explore various foundation models, including Uni-Perceiver v2 [28] (Uni), a Vision Transformer (ViT) pre-trained on LuPerson [41], and the pre-trained image encoder from CLIP [48]. The performance of these diverse foundation models is presented in Tab. 5. As discernible from the table, the performance demonstrates an upward trend with the expansion of the pre-training dataset. Notably, despite LuPerson’s exclusive focus on ReID tasks, its performance lags behind other models due to its comparatively smaller size. This discrepancy underscores the pronounced zero-shot performance benefits

Type	Method	Venue	R→R		I→R		S→R		T→R	
			Rank-1	mAP	Rank-1	mAP	Rank-1	mAP	Rank-1	mAP
Pre-train	LuPerson-NL [12]	CVPR22	24.6*	11.6*	-	-	-	-	-	-
	PLIP [88]	arXiv23	80.4	59.7	-	-	-	-	57.7	-
	APTm [66]	ACMMM23	5.3*	3.5*	-	-	-	-	9.6*	2.7*
Unimodal	OSNet-IBN [84]	ICCV19	73.0	44.9	-	-	-	-	-	-
	M <sup>3</sup> L [81]	CVPR21	78.3	52.5	-	-	-	-	-	-
	OSNet-AIN [85]	TPAMI21	73.3	45.8	-	-	-	-	-	-
Cross-modal	AGW [69]	TPAMI21	17.3*	6.9*	18.2*	19.1*	-	-	-	-
	IRRA [23]	CVPR23	66.6*	40.5*	-	-	-	-	30.1*	25.3*
	UNIReID [7]	CVPR23	19.0*	8.2*	-	-	69.8	73.0	11.6*	9.7*
Multimodal	AIO (Ours)	-	79.6	59.9	57.6	51.9	70.2	73.5	53.4	43.4

Table 7. **Zero-shot performance on cross-modal retrieval.** The best Rank-1 and mAP performance are reported. Results with \* indicate that the experiment results are produced by authors. For AGW [69], it is trained on MSMT17 [58] and LLCM [78] for R→R and I→R. For IRRA [23], it is trained on ICFG-PEDES [9] for T→R. For UNIReID [7], it is trained on Tri-ICFG-PEDES [7] for T→R and R→R.

Method	Rank-1	mAP	mINP
UNIReID (T→R)	76.8	80.6	77.8
UNIReID (S→R)	69.8	73.0	68.3
AIO (T→R)	78.2	81.7	78.4
AIO (S→R)	69.8	72.8	68.8
UNIReID (S+T→R)	91.4	91.8	89.0
AIO (S+T→R)	92.1	92.2	89.2
AIO (R+S+T→R)	93.6	93.7	90.0
AIO (R+I+S+T→R)	93.8	93.7	90.3

Table 8. **Zero-shot performance with multimodal input and generalized cross-modal on PKU-Sketch.**

associated with large-scale pre-trained foundation models.

**Influence of Multimodality Input.** Because our proposed AIO framework supports any combination of diverse modalities as input, we also analyze the influence of different combinations of multimodal inputs on Tri-CUHK-PEDES [7], where the missing IR modality is generated by CA [72]. As presented in Tab. 6, our analysis indicates a preference for RGB and text modalities within our framework over other modalities. Furthermore, when the number of input modalities reaches or exceeds three, there is no significant alteration in performance. This outcome aligns with expectations, as RGB and Text modalities inherently provide more discriminative details than others.

### 4.3. Evaluation on Multimodal ReID

Given the rarity of generalizable works across cross-modal, multimodal, and pre-trained ReID, we conduct a comprehensive comparative analysis involving the proposed AIO framework, various large-scale pre-trained ReID models, unimodal generalized methods, cross-modal methods, and multimodal methods, all within the zero-shot setting. As illustrated in Tab. 7, the existing large-scale pre-trained ReID models, with the exception of PLIP, exhibit unsatisfactory performance in the zero-shot setting. Moreover, AIO achieves competitive performance compared to unimodal generalization methods on R→R retrieval task and outperforms cross-modal methods on all cross-modal re-

trieval tasks in the zero-shot setting. Notably, existing methods fall short in generalizing to unseen modalities, a limitation overcome by AIO, which adeptly handles all four modalities in cross-modal tasks. The outcomes presented in Tab. 8 unveil the remarkable performance of the proposed AIO framework when incorporating multimodal input. This superior performance stands in stark contrast to methods relying solely on unimodal inputs in cross-modal tasks. Additionally, the results consistently underscore the impact of different modalities, aligning with the conclusions drawn from our preceding ablation studies that AIO is more in favor of Text and RGB modalities than others. Moreover, we also discuss the difference between AIO and UNIReID in detail and the limitation of AIO in the supplemental part.

## 5. Conclusion

To the best of our knowledge, this is the **first** work delving into the uncertain multimodal ReID tasks encompassing all four prevalent modalities, *e.g.* RGB, IR, Sketch, and Text. We investigate the feasibility of harnessing large-scale foundation models for multimodal ReID tasks, presenting a prospective avenue toward zero-shot multimodal ReID in wild conditions. In order to cooperate with foundation models, we introduce an innovative multimodal tokenizer, designed to utilize disparate modality inputs within a shared embedding space, guided by carefully crafted cross-modal heads. Moreover, we introduce synthetic augmentation methods with a progressively learning strategy to alleviate the missing modality problem and mitigate the cross-modal gap between different modalities. Extensive experimentation demonstrates the efficacy and competitive performance of the proposed AIO framework across both zero-shot cross-modal and multimodal ReID tasks.

**Acknowledgement.** This work is partially supported by National Natural Science Foundation of China under Grant (62176188, 62361166629, 62225113, 62306215), and the Special Fund of Hubei LuoJia Laboratory (220100015).



## References

- [1] Liping Bao, Longhui Wei, Xiaoyu Qiu, Wengang Zhou, Houqiang Li, and Qi Tian. Learning Transferable Pedestrian Representation from Multimodal Information Supervision, 2023. 1
- [2] Rishi Bommasani and et. al. On the Opportunities and Risks of Foundation Models, 2022. 3
- [3] Adam Botach, Evgenii Zheltonozhskii, and Chaim Baskin. End-to-End Referring Video Object Segmentation With Multimodal Transformers. In *CVPR*, 2022. 3
- [4] Lei Cai, Zhengyang Wang, Hongyang Gao, Dinggang Shen, and Shuiwang Ji. Deep Adversarial Learning for Multi-Modality Missing Data Completion. In *KDD*, 2018. 3
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging Properties in Self-Supervised Vision Transformers. In *ICCV*, 2021. 5
- [6] Cuiqun Chen, Mang Ye, Meibin Qi, and Bo Du. Sketch Transformer: Asymmetrical Disentanglement Learning from Dynamic Synthesis. In *ACMMM*, 2022. 2, 3, 7
- [7] Cuiqun Chen, Mang Ye, and Ding Jiang. Towards Modality-Agnostic Person Re-Identification With Descriptive Query. In *CVPR*, 2023. 1, 2, 3, 6, 7, 8
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *ACL*, 2019. 5
- [9] Zefeng Ding, Changxing Ding, Zhiyin Shao, and Dacheng Tao. Semantically Self-Aligned Network for Text-to-Image Part-aware Person Re-identification, 2021. 8
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *ICLR*, 2021. 5, 7
- [11] Dengpan Fu, Dongdong Chen, Jianmin Bao, Hao Yang, Lu Yuan, Lei Zhang, Houqiang Li, and Dong Chen. Unsupervised Pre-training for Person Re-identification. In *CVPR*, 2021. 2
- [12] Dengpan Fu, Dongdong Chen, Hao Yang, Jianmin Bao, Lu Yuan, Lei Zhang, Houqiang Li, Fang Wen, and Dong Chen. Large-Scale Pre-training for Person Re-identification with Noisy Labels. *CVPR*, 2022. 8
- [13] Takuro Fujii and Shuhei Tarashima. BiLMA: Bidirectional Local-Matching for Text-based Person Re-identification. In *ICCVW*, 2023. 1, 2
- [14] Valentin Gabeur, Chen Sun, Karteek Alahari, and Cordelia Schmid. Multi-modal Transformer for Video Retrieval. In *ECCV*, 2020. 3
- [15] Chenyang Gao, Guanyu Cai, Xinyang Jiang, Feng Zheng, Jun Zhang, Yifei Gong, Pai Peng, Xiaowei Guo, and Xing Sun. Contextual Non-Local Alignment over Full-Scale Representation for Text-Based Person Search, 2021. 3
- [16] Rohit Girdhar, Alaeldin El-Nouby, Zhuang Liu, Mannat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. ImageBind: One Embedding Space To Bind Them All. In *CVPR*, 2023. 3, 7
- [17] Shaojun Gui, Yu Zhu, Xiangxiang Qin, and Xiaofeng Ling. Learning multi-level domain invariant features for sketch re-identification. *Neurocomputing*, 2020. 2, 3
- [18] Shuting He, Hao Luo, Pichao Wang, Fan Wang, Hao Li, and Wei Jiang. TransReID: Transformer-based Object Re-Identification. In *ICCV*, 2021. 3
- [19] Wenke Huang, Mang Ye, Zekun Shi, He Li, and Bo Du. Rethinking Federated Learning With Domain Shift: A Prototype View. In *CVPR*, 2023. 5
- [20] Wenke Huang, Mang Ye, Zekun Shi, Guancheng Wan, He Li, Bo Du, and Qiang Yang. Federated Learning for Generalization, Robustness, Fairness: A Survey and Benchmark, 2023. 3
- [21] Weiran Huang, Mingyang Yi, Xuyang Zhao, and Zihao Jiang. Towards the Generalization of Contrastive Self-Supervised Learning. In *ICLR*, 2023. 1
- [22] Wenke Huang, Yuxia Liu, Mang Ye, Jun Chen, and Bo Du. Federated learning with long-tailed data via representation unification and classifier rectification. 2024. 1
- [23] Ding Jiang and Mang Ye. Cross-Modal Implicit Relation Reasoning and Aligning for Text-to-Image Person Retrieval. In *CVPR*, 2023. 1, 5, 8
- [24] Wonjae Kim, Bokyoung Son, and Ildoo Kim. ViLT: Vision-and-Language Transformer Without Convolution or Region Supervision. In *ICML*, 2021. 3
- [25] Yi-Lun Lee, Yi-Hsuan Tsai, Wei-Chen Chiu, and Chen-Yu Lee. Multimodal Prompting With Missing Modalities for Visual Recognition. In *CVPR*, 2023. 3
- [26] He Li, Mang Ye, and Bo Du. WePerson: Learning a Generalized Re-identification Model from All-weather Virtual Data. In *ACMMM*, 2021. 5
- [27] He Li, Mang Ye, Cong Wang, and Bo Du. Pyramidal Transformer with Conv-Patchify for Person Re-identification. In *ACMMM*, 2022. 3
- [28] Hao Li, Jinguo Zhu, Xiaohu Jiang, Xizhou Zhu, Hongsheng Li, Chun Yuan, Xiaohua Wang, Yu Qiao, Xiaogang Wang, Wenhai Wang, and Jifeng Dai. Uni-Perceiver v2: A Generalist Model for Large-Scale Vision and Vision-Language Tasks. In *CVPR*, 2022. 7
- [29] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before Fuse: Vision and Language Representation Learning with Momentum Distillation. In *NeurIPS*, 2021. 3
- [30] Shuang Li, Tong Xiao, Hongsheng Li, Bolei Zhou, Dayu Yue, and Xiaogang Wang. Person Search With Natural Language Description. In *CVPR*, 2017. 6, 7
- [31] Siyuan Li, Li Sun, and Qingli Li. CLIP-ReID: Exploiting Vision-Language Model for Image Re-Identification without Concrete Text Labels. In *AAAI*, 2023. 3, 7
- [32] Wen Li, Cheng Zou, Meng Wang, Furong Xu, Jianan Zhao, Ruobing Zheng, Yuan Cheng, and Wei Chu. DCFormer: Diverse and Compact Transformer for Person Re-Identification. In *AAAI*, 2023. 2
- [33] Zheng Lian, Lan Chen, Licai Sun, Bin Liu, and Jianhua Tao. GCNet: Graph Completion Network for Incomplete Multimodal Learning in Conversation. *TPAMI*, 2023. 3

- [34] Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. Foundations and Trends in Multimodal Machine Learning: Principles, Challenges, and Open Questions, 2023. 3
- [35] Shengcai Liao and Ling Shao. Interpretable and Generalizable Person Re-Identification with Query-Adaptive Convolution and Temporal Lifting. In *ECCV*, 2020. 5
- [36] Shengcai Liao and Ling Shao. Graph Sampling Based Deep Metric Learning for Generalizable Person Re-Identification. In *CVPR*, 2022. 7
- [37] Kejun Lin, Zhixiang Wang, Zheng Wang, Yinqiang Zheng, and Shin'ichi Satoh. Beyond Domain Gap: Exploiting Subjectivity in Sketch-Based Person Retrieval. In *ACMMM*, 2023. 6, 7
- [38] Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. In *ICLR*, 2019. 7
- [39] Hu Lu, Xuezhong Zou, and Pingping Zhang. Learning Progressive Modality-shared Transformers for Effective Visible-Infrared Person Re-identification. In *AAAI*, 2023. 5
- [40] Hao Luo, Youzhi Gu, Xingyu Liao, Shenqi Lai, and Wei Jiang. Bag of Tricks and A Strong Baseline for Deep Person Re-identification. In *CVPRW*, 2019. 3, 5
- [41] Hao Luo, Pichao Wang, Yi Xu, Feng Ding, Yanxin Zhou, Fan Wang, Hao Li, and Rong Jin. Self-Supervised Pre-Training for Transformer-Based Person Re-Identification. 2021. 2, 5, 7
- [42] Mengmeng Ma, Jian Ren, Long Zhao, Sergey Tulyakov, Cathy Wu, and Xi Peng. SMIL: Multimodal Learning with Severely Missing Modality. In *AAAI*, 2021. 3
- [43] Gary Marcus, Ernest Davis, and Scott Aaronson. A very preliminary analysis of DALL-E 2, 2022. 3
- [44] Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at Once: Enhancing Learning and Generalization Capacities via IBN-Net. In *ECCV*, 2018. 5
- [45] Yongsheng Pan, Mingxia Liu, Yong Xia, and Dinggang Shen. Disease-Image-Specific Learning for Diagnosis-Oriented Neuroimage Synthesis With Incomplete Multi-Modality Data. *TPAMI*, 2021. 3
- [46] Lu Pang, Yaowei Wang, Yi-Zhe Song, Tiejun Huang, and Yonghong Tian. Cross-Domain Adversarial Feature Learning for Sketch Re-identification. In *ACMMM*, 2018. 3, 6, 7
- [47] Petra Poklukar, Miguel Vasco, Hang Yin, Francisco S. Melo, Ana Paiva, and Danica Kragic. Geometric Multimodal Contrastive Representation Learning. In *ICML*, 2022. 2, 3
- [48] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning Transferable Visual Models From Natural Language Supervision. In *ICML*, 2021. 2, 3, 5, 7
- [49] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-Shot Text-to-Image Generation, 2021. 3
- [50] Saikat Roy, Tassilo Wald, Gregor Koehler, Maximilian R. Rokuss, Nico Disch, Julius Holzschuh, David Zimmerer, and Klaus H. Maier-Hein. SAM.MD: Zero-shot medical image segmentation capabilities of the Segment Anything Model, 2023. 3
- [51] Zhiyin Shao, Xinyu Zhang, Meng Fang, Zhifeng Lin, Jian Wang, and Changxing Ding. Learning Granularity-Unified Representations for Text-to-Image Person Re-identification. In *ACMMM*, 2022. 1, 3
- [52] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation Learning with Contrastive Predictive Coding. In *NeurIPS*, 2018. 6
- [53] Patrick von Platen, Suraj Patil, Anton Lozhkov, Pedro Cuenca, Nathan Lambert, Kashif Rasul, Mishig Davaadorj, Dhruv Nair, Sayak Paul, William Berman, Yiyi Xu, Steven Liu, and Thomas Wolf. Diffusers: State-of-the-art diffusion models, 2022. 2, 4, 5, 7
- [54] Andrey Voynov, Kfir Aberman, and Daniel Cohen-Or. Sketch-Guided Text-to-Image Diffusion Models. In *ACM SIGGRAPH*, 2023. 2
- [55] Haochen Wang, Jiayi Shen, Yongtuo Liu, Yan Gao, and Efstratios Gavves. NFormer: Robust Person Re-identification with Neighbor Transformer. *CVPR*, 2022. 2
- [56] Hu Wang, Yuanhong Chen, Congbo Ma, Jodie Avery, Louise Hull, and Gustavo Carneiro. Multi-Modal Learning With Missing Modality via Shared-Specific Feature Modelling. In *CVPR*, 2023. 3
- [57] Rui Wang, Dongdong Chen, Zuxuan Wu, Yinpeng Chen, Xiyang Dai, Mengchen Liu, Yu-Gang Jiang, Luwei Zhou, and Lu Yuan. BEVT: BERT Pretraining of Video Transformers. In *CVPR*, 2022. 2, 3
- [58] Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. Person Transfer GAN to Bridge Domain Gap for Person Re-identification. In *CVPR*, 2018. 8
- [59] Ancong Wu, Wei-Shi Zheng, Shaogang Gong, and Jianhuang Lai. RGB-IR Person Re-identification by Cross-Modality Similarity Preservation. *IJCV*, 2020. 6, 7
- [60] Haiping Wu, Bin Xiao, Noel Codella, Mengchen Liu, Xiyang Dai, Lu Yuan, and Lei Zhang. CvT: Introducing Convolutions to Vision Transformers. In *ICCV*, 2021. 5
- [61] Zesen Wu and Mang Ye. Unsupervised Visible-Infrared Person Re-Identification via Progressive Graph Matching and Alternate Learning. In *CVPR*, 2023. 3, 7
- [62] Bin Yang, Mang Ye, Jun Chen, and Zesen Wu. Augmented Dual-Contrastive Aggregation Learning for Unsupervised Visible-Infrared Person Re-Identification. In *ACMMM*, 2022. 3
- [63] Bin Yang, Jun Chen, Xianzheng Ma, and Mang Ye. Translation, Association and Augmentation: Learning Cross-Modality Re-Identification From Single-Modality Annotation. *TIP*, 2023. 3
- [64] Bin Yang, Jun Chen, and Mang Ye. Towards Grand Unified Representation Learning for Unsupervised Visible-Infrared Person Re-Identification. In *ICCV*, 2023. 1
- [65] Mouxing Yang, Zhenyu Huang, Peng Hu, Taihao Li, Jiancheng Lv, and Xi Peng. Learning With Twin Noisy Labels for Visible-Infrared Person Re-Identification. In *CVPR*, 2022. 2, 3
- [66] Shuyang Yang, Yanan Zhou, Yaxiong Wang, Yujiao Wu, Li Zhu, and Zhedong Zheng. Towards Unified Text-based Person Retrieval: A Large-scale Multi-Attribute and Language Search Benchmark. In *ACMMM*, 2023. 2, 3, 6, 8

- [67] Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. FILIP: Fine-grained Interactive Language-Image Pre-Training. In *ICLR*, 2022. 3
- [68] Mang Ye, Jianbing Shen, David J. Crandall, Ling Shao, and Jiebo Luo. Dynamic Dual-Attentive Aggregation Learning for Visible-Infrared Person Re-Identification. In *ECCV*, 2020. 1, 3
- [69] Mang Ye, Jianbing Shen, Gaojie Lin, Tao Xiang, Ling Shao, and Steven C. H. Hoi. Deep Learning for Person Re-identification: A Survey and Outlook. *TPAMI*, 2021. 1, 2, 3, 7, 8
- [70] Mang Ye, Jianbing Shen, and Ling Shao. Visible-Infrared Person Re-Identification via Homogeneous Augmented Tri-Modal Learning. *TIFS*, 2021. 5
- [71] Mang Ye, Cuiqun Chen, Jianbing Shen, and Ling Shao. Dynamic Tri-Level Relation Mining With Attentive Graph for Visible Infrared Re-Identification. *TIFS*, 2022. 3
- [72] Mang Ye, Zesen Wu, Cuiqun Chen, and Bo Du. Channel Augmentation for Visible-Infrared Re-Identification. *TPAMI*, 2024. 1, 2, 4, 5, 7, 8
- [73] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. CoCa: Contrastive Captioners are Image-Text Foundation Models. 2022. 2
- [74] Yang Yuan. On the Power of Foundation Models. In *ICML*, 2023. 3
- [75] Yajing Zhai, Yawen Zeng, Da Cao, and Shaofei Lu. TriReID: Towards Multi-Modal Person Re-Identification via Descriptive Fusion Model. In *ICMR*, 2022. 2, 3, 6, 7
- [76] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding Conditional Control to Text-to-Image Diffusion Models. In *ICCV*, 2023. 5
- [77] Qiang Zhang, Changzhou Lai, Jianan Liu, Nianchang Huang, and Jungong Han. FMCNet: Feature-Level Modality Compensation for Visible-Infrared Person Re-Identification. In *CVPR*, 2022. 1, 3
- [78] Yukang Zhang and Hanzi Wang. Diverse Embedding Expansion Network and Low-Light Cross-Modality Benchmark for Visible-Infrared Person Re-identification. In *CVPR*, 2023. 6, 7, 8
- [79] Yizhen Zhang, Minkyu Choi, Kuan Han, and Zhongming Liu. Explainable Semantic Space by Grounding Language to Vision with Cross-Modal Contrastive Learning. In *NeurIPS*, 2021. 3
- [80] Yiyuan Zhang, Kaixiong Gong, Kaipeng Zhang, Hongsheng Li, Yu Qiao, Wanli Ouyang, and Xiangyu Yue. Meta-Transformer: A Unified Framework for Multimodal Learning, 2023. 7
- [81] Yuyang Zhao, Zhun Zhong, Fengxiang Yang, Zhiming Luo, Yaojin Lin, Shaozi Li, and Nicu Sebe. Learning to Generalize Unseen Domains via Memory-based Multi-Source Meta-Learning for Person Re-Identification. In *CVPR*, 2021. 8
- [82] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable Person Re-identification: A Benchmark. In *ICCV*, 2015. 6, 7
- [83] Xiangtian Zheng, Xiaohua Huang, Chen Ji, Xiaolin Yang, Pengcheng Sha, and Liang Cheng. Multi-modal person re-identification based on transformer relational regularization. *Information Fusion*, 2024. 5
- [84] Kaiyang Zhou, Yongxin Yang, Andrea Cavallaro, and Tao Xiang. Omni-Scale Feature Learning for Person Re-Identification. In *ICCV*, 2019. 8
- [85] Kaiyang Zhou, Yongxin Yang, Andrea Cavallaro, and Tao Xiang. Learning Generalisable Omni-Scale Representations for Person Re-Identification. *TPAMI*, 2021. 8
- [86] Aichun Zhu, Zijie Wang, Yifeng Li, Xili Wan, Jing Jin, Tian Wang, Fangqiang Hu, and Gang Hua. DSSL: Deep Surroundings-person Separation Learning for Text-based Person Retrieval. In *ACMMM*, 2021. 1, 3
- [87] Zijie Zhuang, Longhui Wei, Lingxi Xie, Tianyu Zhang, Hengheng Zhang, Haozhe Wu, Haizhou Ai, and Qi Tian. Rethinking the Distribution Gap of Person Re-identification with Camera-based Batch Normalization. In *ECCV*, 2020. 5
- [88] Jialong Zuo, Changqian Yu, Nong Sang, and Changxin Gao. PLIP: Language-Image Pre-training for Person Representation Learning, 2023. 6, 7, 8