

IDEA: Inverted Text with Cooperative Deformable Aggregation for Multi-modal Object Re-Identification

Yuhao Wang, Yongfeng Lv, Pingping Zhang*, Huchuan Lu

School of Future Technology, School of Artificial Intelligence, Dalian University of Technology

{924973292, lvyf}@mail.dlut.edu.cn, {zhpp, lhchuan}@dlut.edu.cn

Abstract

Multi-modal object Re-Identification (ReID) aims to retrieve specific objects by utilizing complementary information from various modalities. However, existing methods focus on fusing heterogeneous visual features, neglecting the potential benefits of text-based semantic information. To address this issue, we first construct three text-enhanced multi-modal object ReID benchmarks. To be specific, we propose a standardized multi-modal caption generation pipeline for structured and concise text annotations with Multi-modal Large Language Models (MLLMs). Besides, current methods often directly aggregate multi-modal information without selecting representative local features, leading to redundancy and high complexity. To address the above issues, we introduce IDEA, a novel feature learning framework comprising the Inverted Multi-modal Feature Extractor (IMFE) and Cooperative Deformable Aggregation (CDA). The IMFE utilizes Modal Prefixes and an InverseNet to integrate multi-modal information with semantic guidance from inverted text. The CDA adaptively generates sampling positions, enabling the model to focus on the interplay between global features and discriminative local features. With the constructed benchmarks and the proposed modules, our framework can generate more robust multi-modal features under complex scenarios. Extensive experiments on three multi-modal object ReID benchmarks demonstrate the effectiveness of our proposed method.

1. Introduction

Object Re-Identification (ReID) focuses on retrieving the same object across different camera views. While significant progress has been made with RGB images [9, 13, 23–25, 30, 35, 39, 40, 45, 47, 53], existing methods are limited under environmental challenges like dark and glare, reducing their robustness in real-world applications. Multi-modal imaging, which combines data from multiple spectra such




Figure 1. Overall illustration of our motivations and proposed framework. (a) Our multi-modal caption generation pipeline. (b) Limitations of existing MLLM-generated captions. (c) Comparison between previous methods and our proposed IDEA.

as RGB, Near Infrared (NIR) and Thermal Infrared (TIR), offers a promising solution. By leveraging complementary information, multi-modal ReID methods [4, 26, 37, 54] improve feature robustness under challenging conditions.

Beyond visual information fusion for different spectra, integrating language with vision can enhance ReID performance [11, 27, 34]. However, due to the absence of text annotations for images, most existing methods rely solely on visual features, overlooking the benefits of text-based semantic information. To address this limitation, some studies augment RGB datasets with manual text annotations [6, 20]. While this strategy boosts performance, it is both time-consuming and labor-intensive. With the advent of Multi-modal Large Language Models (MLLMs), image caption has made significant progress. Recently, researchers start to use MLLMs to generate textual descriptions for RGB images [11, 14]. However, these methods face two main challenges. (1) As shown in Fig. 1 (b), the randomness in caption generation often leads to redundancy, resulting in long and complex sentences. Additionally, the text structure varies across images. In ReID tasks, such captions may exceed input length limits of a text encoder [29, 55], with redundant content diluting the essential information. (2) Current methods primarily focus on RGB images, where the generated captions lack sufficient details. However, multi-modal imaging can capture crucial information in complex

*Corresponding author

environments, providing enough information for MLLMs to generate informative captions. **Despite this, no existing methods provide text annotations for multi-modal images.** To bridge this gap, we construct three text-enhanced multi-modal object ReID benchmarks. Meanwhile, another challenge is the aggregation of multi-modal information. As shown in Fig. 1 (c), previous methods directly aggregate the heterogeneous information, leading to high computational complexity and noise interference [54]. To address above issues, we propose a multi-modal feature learning framework named IDEA to introduce Inverted text with cooperative DEformable Aggregation for multi-modal object ReID.

Technically, we first develop a standardized pipeline for multi-modal caption generation. It extends existing datasets with text annotations across spectral modalities. Specifically, as shown in Fig. 1 (a), our caption generation pipeline consists of two steps. In the first step, we use image-template pairs in conjunction with MLLMs to generate informative captions. Then, we leverage the powerful keyword extraction capabilities of MLLMs to extract predefined attributes from the generated captions to produce structured and concise sentences. Building upon the captions, we propose the IDEA framework, which consists of two key components: the Inverted Multi-modal Feature Extractor (IMFE) and Cooperative Deformable Aggregation (CDA). First, we utilize the IMFE to extract integrated multi-modal features using Modal Prefixes and an InverseNet. Specifically, Modal Prefixes are designed to distinguish different modalities, reducing the impact of conflicting information across modalities. The InverseNet further exploits the semantic information from inverted text to enhance feature robustness. Furthermore, we propose CDA to adaptively aggregate discriminative local information. Specifically, based on the aggregated multi-modal information, we adaptively generate sampling positions to enable the model to focus on the interplay between global features and discriminative local information. Through these modules, our proposed framework effectively utilizes semantic guidance from texts while adaptively aggregating discriminative multi-modal information. Extensive experiments on three benchmark datasets demonstrate the effectiveness of our approach. In summary, our contributions are as follows:

- We construct three text-enhanced multi-modal object ReID benchmarks, providing a structured caption generation pipeline across multiple spectral modalities.
- We introduce a novel feature learning framework named IDEA, which includes the Inverted Multi-modal Feature Extractor (IMFE) and Cooperative Deformable Aggregation (CDA). The IMFE integrates multi-modal features using Modal Prefixes and an InverseNet, while the CDA adaptively aggregates discriminative local information.
- Extensive experiments on three benchmark datasets validate the effectiveness of our proposed method.

2. Related Work

2.1. Vision-language Learning in Re-Identification

Vision-language learning plays a pivotal role in multi-modal applications. As the field advances [5, 48, 49], there is a growing focus on exploring the interaction between visual and textual information in ReID tasks. Existing methods can be broadly divided into cross-modal and multi-modal ReID. In cross-modal ReID, Text-to-Image ReID [6] focuses on matching text queries with image galleries. Current methods primarily emphasize feature alignment [17] and pre-training [32]. To expand application scenarios, Zhai et al. [52] combine text and sketch modalities in the query, while Chen et al. [3] propose an uncertain Query-to-RGB retrieval model to handle missing modalities. More recently, Li et al. [19] introduce flexible modality combinations and unify various uncertainties in retrieval. However, these methods focus on cross-modal alignment. In contrast, multi-modal ReID aims to leverage complementary information from multiple modalities, with both the query and gallery containing multi-modal data. The rise of vision-language models, particularly CLIP [29], has significantly advanced this area by facilitating image-text interactions. Li et al. [21] pioneer CLIP-ReID to learn from text templates for ReID tasks. However, the templates are not real descriptions of images, potentially restricting performance. To address this, Han et al. [11] leverage MLLMs to generate descriptions for RGB images. Despite improvements, existing methods fail to tackle the challenges of generating text for multi-spectral modalities and ensuring structural consistency in the generated descriptions. To address the above issues, we propose a structured multi-modal caption generation pipeline, enhancing the consistency of texts and providing informative annotations for multi-modal object ReID.

2.2. Multi-modal Object Re-Identification

Multi-modal object ReID gains great attention due to its robustness in real-world applications. Research primarily focuses on leveraging complementary information from different modalities. Early works emphasize effective fusion strategies [8, 56]. To enhance modality-specific learning, Wang et al. [41] propose an interact-embed-enlarge framework to facilitate knowledge expansion. To address missing modal data, Zheng et al. [57] introduce pixel reconstruction to handle data inconsistency. Li et al. [18] use a coherence loss to guide feature fusion. Some methods further improve feature robustness with modality generation, graph learning and instance sampling [10, 12, 58]. Notably, Wang et al. [42] introduce a test-time training framework to mine inter-modal interactions. Wang et al. [37] develop TOP-ReID, which employs token permutation to guide modality fusion. Recently, Zhang et al. [54] propose diverse token selections to mitigate background noises. Despite promis-




Figure 2. The upper row presents the template used to annotate our dataset. The lower row provides detailed information about the annotated dataset. (a) Category statistics for our annotated person ReID dataset. (b) Example images and captions from the RGBNT201 dataset. (c) Category statistics for our annotated vehicle ReID dataset. (d) Example images and captions from the MSVR310 dataset.

ing results, they often overlook the semantic guidance of informative text annotations. Thus, we propose the IDEA framework, which incorporates generated text annotations to leverage semantic guidance. This framework adaptively aggregates discriminative local information, enhancing feature robustness in complex scenarios.

3. Proposed Method

In this section, we introduce the multi-modal caption generation pipeline and the proposed feature learning framework. Fig. 2 illustrates the details of caption generation, including attribute definitions and examples. Fig. 3 presents the key modules of our proposed IDEA: the Inverted Multi-modal Feature Extractor (IMFE) and the Cooperative Deformable Aggregation (CDA). Details are described as follows.

3.1. Multi-modal Caption Generation

To bridge the gap in multi-modal text annotation, we propose the Multi-modal Caption Generation pipeline. To be specific, our pipeline consists of two steps. The first step generates informative descriptions, while the second step reuses the MLLMs to extract the predefined key attributes.

Caption Generation. Taking our person annotation as an example, we define 8 attribute categories in Fig. 2 (a), including gender, belongings and so on. These attributes are mapped into the template shown in the top part of Fig. 2 (a). To tailor the description to each modality, we add specific prefixes: "Write a comprehensive description of the person's overall appearance based on the [RGB/NIR/TIR] image, ...". Besides, we incorporate commands to prevent the hallucination of the MLLMs [32], forming a complete prompt. This prompt, along with the corresponding image, is fed into the MLLMs to generate the textual description.

Attribute Extraction. After obtaining an initial descrip-

tion, we feed it back into the MLLMs to extract predefined attributes. After that, we populate these attributes into the same template we used for caption generation, forming a structured and concise description. As shown in Fig. 2 (b), the descriptions for each modality can align well with the corresponding images. For instance, the "white t-shirt" and "purse" in the RGB image are accurately captured. For the vehicle dataset, we follow similar pipelines. As illustrated in Fig. 2 (c), the text of the RGB image successfully captures the license plate number "KQQ819". Meanwhile, the car logo is more prominent in the TIR image, with "Hyundai" being identified. This approach allows us to establish a unified pipeline for multi-modal text annotation, ensuring structural descriptions across different modalities.

3.2. IDEA Framework

3.2.1. Inverted Multi-modal Feature Extractor

To leverage the semantic guidance from texts, we propose the Inverted Multi-modal Feature Extractor (IMFE). Unlike previous methods [37, 54] that focus on fusing multi-modal images, we incorporate multi-modal texts into the fusion. However, as indicated by the blue keywords in Fig. 2 (b) and (c), conflicts can arise due to the nature of multi-modal images. Directly aggregating contradictory information can lead to model confusion. Thus, IMFE employs Modal Prefixes and an InverseNet to address these issues.

Modal Prefixes. To enhance the model's awareness of different modal texts and mitigate the impact of conflicting information, we propose a simple method called Modal Prefixes. Taking the text input for the RGB branch as an example, as shown in the top left corner of Fig. 3, T_{RGB} represents the text annotation of the RGB image, while P_{RGB} is the modal prefix for RGB modality. P_{RGB} consists of two parts: a fixed text describing the characteristics of the




Figure 3. Illustration of the proposed IDEA framework. The upper part depicts the Inverted Multi-modal Feature Extractor (IMFE). It employs modal prefixes and an InverseNet to incorporate semantic text guidance for feature discriminability. The lower part highlights the Cooperative Deformable Aggregation (CDA), which adaptively integrates discriminative local information with global features. With the integration of IMFE and CDA, IDEA effectively extracts discriminative multi-modal features for object ReID.

RGB image and learnable tokens for fine-tuning. Specifically, P_{RGB} is structured as: “An image of a XXXX person in the visible spectrum, capturing natural colors and fine details: ”, where “XXXX” is replaced by an equal number of learnable tokens [22]. Here, we denote the number of learnable tokens as N_P . Then, we concatenate P_{RGB} and T_{RGB} to form the text input T_R for the RGB branch:

$$T_R = [P_{RGB}, T_{RGB}]. \quad (1)$$

Other modalities follow the same structure to guide the model in distinguishing between different modal texts. With this simple yet effective method, the model can better understand the characteristics of each modality and reduce conflicts during multi-modal fusion.

InverseNet. To fully leverage the semantic information from texts, we propose the InverseNet. Previous methods [2, 11] often reverse global image information into pseudo-text tokens and use texts as the primary modality for downstream tasks. However, due to the limitations of MLLMs and the complexity of multi-modal image annotations, the generated texts may contain some errors and conflicts. To address this, we take a different direction, reversing global text information into pseudo-image tokens. In this reversal, we also incorporate rich information from the learnable tokens in Modal Prefixes. This approach enables the interaction between semantic information and image details, resulting in a discriminative feature representation.

Specifically, as shown in the upper part of Fig. 3, the text T_m is first fed into the CLIP’s text encoder $\mathcal{T}$ for extracting the semantic feature $\hat{f}_m^t \in \mathbb{R}^C$ as follows:

$$\hat{f}_m^t = \mathcal{T}(T_m). \quad (2)$$

Here, $m \in \{R, N, T\}$ represents the RGB, NIR and TIR modalities, respectively. C denotes the embedding dimension. If learnable tokens exist in the modal prefixes, $\hat{f}_m^t$ will be updated by the average of learnable tokens. Then, we feed $\hat{f}_m^t$ into the InverseNet $\mathcal{I}$ to generate the pseudo token $f_m^t \in \mathbb{R}^{C'}$ with the following equation:

$$f_m^t = \mathcal{I}(\hat{f}_m^t). \quad (3)$$

Here, C' is the dimension of the pseudo token **while $\mathcal{I}$ is essentially a simple layer of MLP**. Meanwhile, for the image input I_m , we first patchify it into N_l patches. Then, we concatenate the pseudo token f_m^t , the patch tokens and a learnable token f_m^v to form the image feature $\bar{F}_m \in \mathbb{R}^{(N_l+2) \times C'}$. Finally, we feed $\bar{F}_m$ into the CLIP’s vision encoder $\mathcal{V}$ to obtain the integrated feature $F_m \in \mathbb{R}^{(N_l+2) \times C}$ as follows:

$$F_m = \mathcal{V}(\bar{F}_m). \quad (4)$$

By incorporating the InverseNet, we can effectively leverage the semantic information from texts and enhance the discriminability of the feature representation.

3.2.2. Cooperative Deformable Aggregation

To effectively extract discriminative local information and adaptively aggregate complementary multi-modal features, we propose Cooperative Deformable Aggregation (CDA). CDA addresses the noise interference by adaptively generating sampling positions of key local regions. This module enables interactions between global features and discriminative local features, ensuring efficient feature aggregation.

Technically, as illustrated in the lower part of Fig. 3, we first extract patch tokens from the integrated feature F_m .

These patches are then reshaped to form $\hat{F}_m \in \mathbb{R}^{H \times W \times C}$, where H and W represent the height and width of the local feature map, respectively. Next, we concatenate these features along the channel dimension and pass them into the Local Mixer $\mathcal{M}$ for interaction. Specifically, $\mathcal{M}$ comprises point-wise convolution $\mathcal{P}$, the GELU activation function [15] δ and depth-wise convolution $\mathcal{D}$. This process yields the aggregated local features $\hat{F}_A \in \mathbb{R}^{H_S \times W_S \times C}$:

$$\hat{F}_A = \mathcal{M}([\hat{F}_R, \hat{F}_N, \hat{F}_T]), \quad (5)$$

$$\mathcal{M}(\mathcal{X}) = \delta(\mathcal{D}(\delta(\mathcal{P}(\mathcal{X})))), \quad (6)$$

where $[\cdot]$ means concatenation. H_S and W_S represent the number of sampled regions in the height and width dimensions, respectively. Since the convolution $\mathcal{D}$ aggregates local information in the receptive field of the kernel, we define the center of each convolution operation as the reference points $P \in \mathbb{R}^{H_S \times W_S \times 2}$. Based on these reference points, the aggregated local information is then used to determine the direction in which the model should shift within the local region. Thus, we send $\hat{F}_A$ into the Offset Generator $\mathcal{G}$ to predict the offsets $\Delta P \in \mathbb{R}^{H_S \times W_S \times 2}$ as follows:

$$\Delta P = \mathcal{G}(\hat{F}_A) * k, \quad (7)$$

where k scales the offset magnitude [44] and $\mathcal{G}$ denotes a linear layer. Finally, we apply the offsets ΔP to the reference points P to get the sampling locations $\hat{P} \in \mathbb{R}^{H_S \times W_S \times 2}$, which are **shared across modalities** [54]:

$$\hat{P} = P + \Delta P. \quad (8)$$

Next, we sample the local features $\hat{F}_m$ at the positions $\hat{P}$ using the bilinear interpolation to obtain the discriminative local features. The features are concatenated to form the sampled feature $F_S \in \mathbb{R}^{3N_S \times C}$, where $N_S = H_S \times W_S$ represents the number of sampled regions. At this stage, we utilize the cooperative information from multiple modalities to generate deformable local discriminative features. As illustrated by the sampling points visualization in Fig. 3, the adaptive offsets effectively concentrate on critical semantic regions of the human body, enhancing the discrimination.

To promote the interaction between global features and discriminative local features, we incorporate cross attention [7]. Specifically, we extract the visual class token and pseudo token from the integrated feature F_m to construct the global feature $F_G \in \mathbb{R}^{6 \times C}$. Then, the global feature F_G is used as the query, while the sampled local feature F_S serves as the key and value to facilitate interaction:

$$F_{CDA} = F_G + \mathcal{CA}(F_G, F_S), \quad (9)$$

where $F_{CDA} \in \mathbb{R}^{6 \times C}$ represents the final feature while $\mathcal{CA}$ denotes the cross attention block. By incorporating the CDA, we can effectively extract discriminative local information and adaptively aggregate complementary multi-modal features, enhancing the model’s discriminability.

	Methods	RGBNT201			
		mAP	R-1	R-5	R-10
Multi-modal	HAMNet [18]	27.7	26.3	41.5	51.7
	PFNet [56]	38.5	38.9	52.0	58.4
	IEEE [41]	47.5	44.4	57.1	63.6
	DENet [57]	42.4	42.2	55.3	64.5
	LRMM [43]	52.3	53.4	64.6	73.2
	UniCat* [4]	57.0	55.7	-	-
	HTT* [42]	71.1	73.4	83.1	87.3
	TOP-ReID* [37]	72.3	76.6	84.7	89.4
	EDITOR* [54]	66.5	68.3	81.1	88.2
	RSCNet* [50]	68.2	72.5	-	-
	WTSF-ReID* [51]	67.9	72.2	83.4	89.7
	MambaPro† [36]	78.9	83.4	<u>89.8</u>	91.9
	DeMo† [38]	<u>79.0</u>	<u>82.3</u>	88.8	<u>92.0</u>
	IDEA†	80.2	82.1	90.0	93.3

Table 1. Performance comparison on RGBNT201. Best results are in bold, the second bests are underlined. † denotes CLIP-based methods, * indicates ViT-based while others are CNN-based ones.

3.3. Objective Function

As shown in Fig. 3, we optimize IDEA through losses applied to multiple features. To maintain consistency with prior works [37, 54], we separately extract image and text features from F_G and F_{CDA} , resulting in F_G^v, F_G^t and $F_{CDA}^v, F_{CDA}^t \in \mathbb{R}^{3C}$. For each feature, we apply label smoothing cross-entropy loss [31] and triplet loss [16]:

$$\mathcal{L}_g(\mathcal{F}) = \mathcal{L}_{CE}(\mathcal{F}) + \mathcal{L}_{Tri}(\mathcal{F}), \quad (10)$$

where $\mathcal{F}$ denotes the input feature. Here, $\mathcal{L}_{CE}$ is the label smoothing cross-entropy loss while $\mathcal{L}_{Tri}$ means the triplet loss. The overall objective function is then formulated as:

$$\mathcal{L} = \mathcal{L}_g(F_G^v) + \mathcal{L}_g(F_{CDA}^v) + \mathcal{L}_g(F_G^t) + \mathcal{L}_g(F_{CDA}^t). \quad (11)$$

4. Experiments

4.1. Datasets and Evaluation Protocols

Datasets. We evaluate the proposed method on three multi-modal object ReID benchmarks. To efficiently annotate multi-spectral images without hardware constraints, we employ the API-based Qwen-VL [1] for automated textual description generation. To be specific, RGBNT201 [56] comprises 4,787 image triplets with 14,361 annotations, each averaging 35.47 characters and encompassing 8 distinct attributes. MSVR310 [58] contains 2,087 image triplets and 6,261 annotations, with an average length of 56.06 characters and 6 attributes. RGBNT100 [18] includes the largest number of triplets at 17,250 and 51,750 annotations with an average length of 56.25 characters and 6 attributes.

Evaluation Protocols. The performance is evaluated using

	Methods	RGBNT100		MSVR310	
		mAP	R-1	mAP	R-1
Multi-modal	HAMNet [18]	74.5	93.3	27.1	42.3
	PFNet [56]	68.1	94.1	23.5	37.4
	GAFNet [10]	74.4	93.4	-	-
	GPFNet [12]	75.0	94.5	-	-
	CCNet [58]	77.2	96.3	36.4	55.2
	LRMM [43]	78.6	96.7	36.7	49.7
	GraFT* [46]	76.6	94.3	-	-
	UniCat* [4]	79.4	96.2	-	-
	PHT* [28]	79.9	92.7	-	-
	HTT* [42]	75.7	92.6	-	-
	TOP-ReID* [37]	81.2	96.4	35.9	44.6
	EDITOR* [54]	82.1	96.4	39.0	49.3
	FACENet* [59]	81.5	96.9	36.2	54.1
	RSCNet* [50]	82.3	96.6	39.5	49.6
	WTSF-ReID* [51]	82.2	96.5	39.2	49.1
	MambaPro [†] [36]	83.9	94.7	47.0	56.5
	DeMo [†] [38]	86.2	97.6	49.2	59.8
	IDEA[†]	87.2	96.5	47.0	62.4

Table 2. Performance on RGBNT100 and MSVR310.

mean Average Precision (mAP) and Cumulative Matching Characteristics (CMC) at Rank-K ($K = 1, 5, 10$).

4.2. Implementation Details

Our model is implemented using PyTorch and trained on an NVIDIA A800 GPU. The pre-trained model CLIP [29] is used for both the vision and text encoders. For the datasets, images triples are resized to 256×128 for RGBNT201 and 128×256 for MSVR310 and RGBNT100. Data augmentation includes random horizontal flipping, cropping and random erasing [60]. For RGBNT201 and MSVR310, the mini-batch size is set to 64, with 8 images sampled per identity. For RGBNT100, the mini-batch size is set to 128, with 16 images sampled per identity. We fine-tune the learnable modules using an initial learning rate of 3.5×10^{-6} , which decays to 3.5×10^{-7} during training. For the RGBNT201 dataset, the values of N_p and k are set to 2 and 5, respectively. Other details are provided in the supplementary material. **Code and annotations are available at IDEA.**

4.3. Comparison with State-of-the-Art Methods

Multi-modal Person ReID. In Tab. 1, we compare our method IDEA[†] with existing multi-modal approaches on the RGBNT201 dataset. Leveraging complementary information from different modalities, multi-modal methods exhibit superior performance. Specifically, our proposed IDEA[†] achieves 80.2% mAP and 82.1% Rank-1 accuracy, surpassing TOP-ReID* by 7.9% in mAP and 5.5% in Rank-1. Compared with other leading methods like HTT* and

Index	Modules			Metrics	
	Text	IMFE	CDA	mAP	Rank-1
A	×	×	×	70.3	72.1
B	✓	×	×	73.4	75.8
C	✓	✓	×	77.2	81.1
D	✓	✓	✓	80.2	82.1

Table 3. Performance comparison with different modules.

Index	IMFE			Metrics	
	InverseNet	Prefixes	Prompt	mAP	Rank-1
A	×	×	-	72.6	75.1
B	×	✓	×	73.4	75.8
C	✓	×	-	73.7	77.3
D	✓	✓	×	75.4	78.6
E	✓	✓	✓	77.2	81.1

Table 4. Comparison with different components in IMFE.

EDITOR*, IDEA[†] shows a superior adaptability to complex scenarios, confirming the robustness. Besides, IDEA[†] outperforms the CLIP-based methods MambaPro[†] and DeMo[†] by 1.3% and 1.2% in mAP, respectively. These results highlight IDEA’s ability to utilize semantic information from textual guidance for improved feature discrimination.

Multi-modal Vehicle ReID. In Tab. 2, we compare IDEA[†] with state-of-the-art methods on the RGBNT100 and MSVR310 datasets. Our IDEA[†] achieves an mAP of 87.2%, surpassing EDITOR* by 5.1% in mAP. Especially on the challenging MSVR310 dataset, IDEA[†] achieves an mAP of 47.0% and a Rank-1 accuracy of 62.4%, outperforming EDITOR* by 8.0% in mAP and 13.1% in Rank-1 accuracy. These results verify IDEA’s generalization ability.

4.4. Ablation Studies

We evaluate the effectiveness of the proposed modules on the RGBNT201 dataset. To be specific, our baseline comprises a three-branch vision encoder, which performs retrieval by concatenating the class tokens from three modalities. Upon incorporating the IMFE, we utilize F_G^t for retrieval and when the CDA is included, F_{CDA}^t is employed.

Effects of Key Modules. Tab. 3 presents the performance of various combinations of our proposed modules. Model A serves as the baseline, achieving an mAP of 70.3% and Rank-1 accuracy of 72.1%. Model B incorporates text information through parallel text encoders, improving the mAP to 73.4% and Rank-1 accuracy to 75.8%. With the inverted structure in IMFE, Model C further enhances performance, reaching an mAP of 77.2% and Rank-1 of 81.1%. Finally, Model D integrates CDA, delivering the best performance with an mAP of 80.2%. These results fully validate the effectiveness and efficiency of the proposed modules.

Effects of Key Components in IMFE. Tab. 4 shows the performance of different components in the IMFE. Model A represents the parallel text encoders without the inverted structure or prefixes, achieving an mAP of 72.6% and Rank-

Index	CDA			Metrics	
	Sample	Cross Attn	Shared Offset	mAP	Rank-1
A	×	×	-	76.3	78.7
B	×	✓	-	77.0	79.8
C	✓	×	×	76.8	78.9
D	✓	×	✓	77.6	80.4
E	✓	✓	×	79.5	81.7
F	✓	✓	✓	80.2	82.1

Table 5. Comparison with different components in CDA.

Model	mAP	Rank-1	Rank-5	Rank-10
IDEA w/o Text	74.5	75.0	84.8	88.8
IDEA	80.2	82.1	90.0	93.3

Table 6. Comparison of IDEA with and without text.

Model	mAP	Rank-1	Rank-5	Rank-10
IDEA w/o Offset	78.4	81.3	89.7	92.2
IDEA	80.2	82.1	90.0	93.3

Table 7. Comparison of IDEA with and without offset.

1 accuracy of 75.1%. With the Modal Prefixes in Model B, the mAP increases to 73.4%, demonstrating the effectiveness of our prefixes mechanism. Model C introduces the inverted structure, achieving an mAP of 73.7%. Model D combines the inverted structure and prefixes, further improving the mAP to 75.4%. Finally, Model E integrates the learnable tokens in Modal Prefixes, achieving the best performance with an mAP of 77.2%. Overall, inverted structure performs better than the parallel structure, achieving better performance. These results demonstrate the effectiveness of different components in our proposed IMFE.

Effects of Key Components in CDA. Based on the IMFE, Tab. 5 shows the performance of various components within CDA. When sampling is excluded, all local patches in $\hat{F}_m$ are processed. Without cross attention, multi-modal interaction uses only local features, omitting global context, with pooled $\hat{F}_m$ for retrieval. If shared offset is excluded, each modality independently generates offsets. Regarding numerical results, Model A, which does not incorporate these components, achieves an mAP of 76.3% and Rank-1 accuracy of 78.7%. Model B introduces cross attention, leading to a slight improvement in mAP to 77.0%. Model C adds sampling, resulting in a 0.5% improvement in mAP over Model A by reducing noise in the local features. Model D incorporates shared offset, yielding a 0.8% improvement in mAP and a 1.5% increase in Rank-1 accuracy compared to Model C. Model E combines cross attention and sampling, achieving an mAP of 79.5%. Finally, Model F delivers the best performance with an mAP of 80.2%. These results validate the effectiveness of each component within CDA.

Effects of Text Guidance in IDEA. Tab. 6 compares IDEA’s performance with and without text input. Model A, lacking text input, achieves an mAP of 74.5%. With text guidance, Model B improves to 80.2% mAP, demonstrating the effectiveness of text in enhancing feature robustness.




Figure 4. Comparison with different hyper-parameters.




Figure 5. Visualization of the cosine similarity distribution.

Effects of Offset in IDEA. Tab. 7 compares the performance of IDEA with and without the offset mechanism. Model A uses aggregated local features with convolution to interact with global features, achieving an mAP of 78.4% and Rank-1 accuracy of 81.3%. Model B introduces the offset mechanism, improving the mAP to 80.2% and Rank-1 accuracy to 82.1%. This demonstrates that the offset mechanism helps the model focus on more important regions.

Effects of Prompt Lengths and Offset Factors. Fig. 4 illustrates the performance of IDEA with varying prompt lengths and offset factors. The results show that a prompt length of 2 and an offset factor of 5 yield the best performance. As the prompt length increases, the model’s performance improves, but the gains become smaller as the prompt length grows, likely due to the text truncation. Regarding the offset factor, the model remains stable.

4.5. Visualization Analysis

Cosine Similarity Distributions. Fig. 5 presents the distributions of cosine similarities among test features. The results verify that incorporating text enhances feature discrimination, while adding offset mechanism further amplifies the distinction between positive and negative samples, confirming the capture of more discriminative features.

Multi-modal Feature Distributions. Fig. 6 visualizes the feature distributions of different modules. Comparing Fig. 6 (a) and (b), the introduction with text guidance improves feature discrimination, as instances of the same ID become more compact. In Fig. 6 (c), with IMFE, the feature distributions become more discriminative than the parallel structure in Fig. 6 (b). Finally, in Fig. 6 (d), CDA further enhances feature discrimination. These visualizations demonstrate the effectiveness of our proposed modules.




Figure 6. Visualization of the feature distributions with t-SNE [33]. Different colors represent different identities.




Figure 7. Visualization of the generated offsets.



Figure 8. Rank list comparison between the baseline and IDEA.

Visualization of the Generated Offsets. Fig. 7 visualizes the offsets generated by CDA. Each arrow starts at the reference point and ends at the feature sampling point. We map these offsets onto the original image to identify the regions where the model shifts its attention. As shown, the offsets highlight discriminative areas, such as the head, bag and shoes. It demonstrates how the offset mechanism guides the model to focus on the most discriminative regions, effectively reducing the influence of noisy information.

Rank List Comparison. Fig. 8 compares the cross-camera rank lists from the baseline and IDEA. IDEA produces more accurate rankings, while the baseline yields noisier results. These visualizations validate the effectiveness of IDEA.




Figure 9. Visualization of the channel activation maps. Each set includes the original image, baseline and IDEA map, respectively.

Visualization of Channel Activation Maps. Fig. 9 compares the channel activation maps of our baseline model and IDEA. The incorporation of textual guidance helps the model focus on more discriminative regions, enhancing feature robustness and improving interpretability.

5. Conclusion

In this work, we propose IDEA, a novel feature learning framework for multi-modal object ReID. We first construct three text-enhanced multi-modal object ReID benchmarks using MLLMs, providing a structured caption generation pipeline. With the generated text, the Inverted Multi-modal Feature Extractor (IMFE) leverages semantic guidance from inverted texts while reducing fusion conflicts. Additionally, the Cooperative Deformable Aggregation (CDA) adaptively aggregates discriminative local features with global information. Experiments on three public ReID benchmarks demonstrate the effectiveness of our method.

Acknowledgements. This work was supported in part by the National Natural Science Foundation of China (No. 62101092, 62476044, 62388101) and Open Project of Anhui Provincial Key Laboratory of Multimodal Cognitive Computation, Anhui University (No. MMC202102, MMC202407).

References

- [1] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*, 1(2):3, 2023. 5
- [2] Alberto Baldrati, Lorenzo Agnolucci, Marco Bertini, and Alberto Del Bimbo. Zero-shot composed image retrieval with textual inversion. In *ICCV*, pages 15338–15347, 2023. 4
- [3] Cuiqun Chen, Mang Ye, and Ding Jiang. Towards modality-agnostic person re-identification with descriptive query. In *CVPR*, pages 15128–15137, 2023. 2
- [4] Jennifer Crawford, Haoli Yin, Luke McDermott, and Daniel Cummings. Unicat: Crafting a stronger fusion baseline for multimodal re-identification. *arXiv preprint arXiv:2310.18812*, 2023. 1, 5, 6
- [5] Haiwen Diao, Yufeng Cui, Xiaotong Li, Yuezhe Wang, Huchuan Lu, and Xinlong Wang. Unveiling encoder-free vision-language models. *arXiv preprint arXiv:2406.11832*, 2024. 2
- [6] Zefeng Ding, Changxing Ding, Zhiyin Shao, and Dacheng Tao. Semantically self-aligned network for text-to-image part-aware person re-identification. *arXiv preprint arXiv:2107.12666*, 2021. 1, 2
- [7] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 5
- [8] Yunpeng Gong, Liqing Huang, and Lifei Chen. Eliminate deviation with deviation for data augmentation and a general multi-modal data learning method. *arXiv preprint arXiv:2101.08533*, 2021. 2
- [9] Yunpeng Gong, Zhun Zhong, Yansong Qu, Zhiming Luo, Rongrong Ji, and Min Jiang. Cross-modality perturbation synergy attack for person re-identification. In *NeurIPS*, 2024. 1
- [10] Jinbo Guo, Xiaojing Zhang, Zhengyi Liu, and Yuan Wang. Generative and attentive fusion for multi-spectral vehicle re-identification. In *ICSP*, pages 1565–1572, 2022. 2, 6, 5
- [11] Qianru Han, Xinwei He, Zhi Liu, Sannyuya Liu, Ying Zhang, and Jinhai Xiang. Clip-scgi: Synthesized caption-guided inversion for person re-identification. *arXiv preprint arXiv:2410.09382*, 2024. 1, 2, 4
- [12] Qiaolin He, Zefeng Lu, Zihan Wang, and Haifeng Hu. Graph-based progressive fusion network for multi-modality vehicle re-identification. *TITS*, pages 1–17, 2023. 2, 6
- [13] Shuting He, Hao Luo, Pichao Wang, Fan Wang, Hao Li, and Wei Jiang. Transreid: Transformer-based object re-identification. In *ICCV*, pages 15013–15022, 2021. 1
- [14] Weizhen He, Yiheng Deng, Shixiang Tang, Qihao Chen, Qingsong Xie, Yizhou Wang, Lei Bai, Feng Zhu, Rui Zhao, Wanli Ouyang, et al. Instruct-reid: A multi-purpose person re-identification task with instructions. In *CVPR*, pages 17521–17531, 2024. 1
- [15] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016. 5
- [16] Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*, 2017. 5
- [17] Ding Jiang and Mang Ye. Cross-modal implicit relation reasoning and aligning for text-to-image person retrieval. In *CVPR*, pages 2787–2797, 2023. 2
- [18] Hongchao Li, Chenglong Li, Xianpeng Zhu, Aihua Zheng, and Bin Luo. Multi-spectral vehicle re-identification: A challenge. In *AAAI*, pages 11345–11353, 2020. 2, 5, 6
- [19] He Li, Mang Ye, Ming Zhang, and Bo Du. All in one framework for multimodal re-identification in the wild. In *CVPR*, pages 17459–17469, 2024. 2
- [20] Shuang Li, Tong Xiao, Hongsheng Li, Bolei Zhou, Dayu Yue, and Xiaogang Wang. Person search with natural language description. In *CVPR*, pages 1970–1979, 2017. 1
- [21] Siyuan Li, Li Sun, and Qingli Li. Clip-reid: exploiting vision-language model for image re-identification without concrete text labels. In *AAAI*, pages 1405–1413, 2023. 2
- [22] Yaowei Li, Zimo Liu, Wenming Yang, Yaowei Wang, Qingmin Liao, et al. Clip-based synergistic knowledge transfer for text-based person retrieval. *arXiv preprint arXiv:2309.09496*, 2023. 4
- [23] Xuehu Liu, Pingping Zhang, Chenyang Yu, Huchuan Lu, and Xiaoyun Yang. Watching you: Global-guided reciprocal learning for video-based person re-identification. In *CVPR*, pages 13334–13343, 2021. 1
- [24] Xuehu Liu, Chenyang Yu, Pingping Zhang, and Huchuan Lu. Deeply coupled convolution–transformer with spatial–temporal complementary learning for video-based person re-identification. *TNNLS*, 2023.
- [25] Xuehu Liu, Pingping Zhang, Chenyang Yu, Xuesheng Qian, Xiaoyun Yang, and Huchuan Lu. A video is worth three views: Trigeminal transformers for video-based person re-identification. *TITS*, 2024. 1
- [26] Hu Lu, Xuezhong Zou, and Pingping Zhang. Learning progressive modality-shared transformers for effective visible-infrared person re-identification. In *AAAI*, pages 1835–1843, 2023. 1
- [27] Ke Niu, Haiyang Yu, Mengyang Zhao, Teng Fu, Siyang Yi, Wei Lu, Bin Li, Xuelin Qian, and Xiangyang Xue. Chatreid: Open-ended interactive person retrieval via hierarchical progressive tuning for vision language models. *arXiv preprint arXiv:2502.19958*, 2025. 1
- [28] Wenjie Pan, Linhan Huang, Jianbao Liang, Lan Hong, and Jianqing Zhu. Progressively hybrid transformer for multi-modal vehicle re-identification. *Sensors*, 23(9):4206, 2023. 6
- [29] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021. 1, 2, 6
- [30] Jiangming Shi, Xiangbo Yin, Yachao Zhang, Yuan Xie, Yanyun Qu, et al. Learning commonality, divergence and variety for unsupervised visible-infrared person re-identification. *NeurIPS*, 37:99715–99734, 2024. 1

- [31] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, pages 2818–2826, 2016. 5
- [32] Wentan Tan, Changxing Ding, Jiayu Jiang, Fei Wang, Yibing Zhan, and Dapeng Tao. Harnessing the power of mllms for transferable text-to-image person reid. In *CVPR*, pages 17127–17137, 2024. 2, 3, 1
- [33] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *JMLR*, 9(11), 2008. 8
- [34] Qizao Wang, Bin Li, and Xiangyang Xue. When large vision-language models meet person re-identification. *arXiv preprint arXiv:2411.18111*, 2024. 1
- [35] Yingquan Wang, Pingping Zhang, Shang Gao, Xia Geng, Hu Lu, and Dong Wang. Pyramid spatial-temporal aggregation for video-based person re-identification. In *ICCV*, pages 12026–12035, 2021. 1
- [36] Yuhao Wang, Xuehu Liu, Tianyu Yan, Yang Liu, Aihua Zheng, Pingping Zhang, and Huchuan Lu. Mamba-pro: Multi-modal object re-identification with mamba aggregation and synergistic prompt. *arXiv preprint arXiv:2412.10707*, 2024. 5, 6
- [37] Yuhao Wang, Xuehu Liu, Pingping Zhang, Hu Lu, Zhengzheng Tu, and Huchuan Lu. Top-reid: Multi-spectral object re-identification with token permutation. In *AAAI*, pages 5758–5766, 2024. 1, 2, 3, 5, 6
- [38] Yuhao Wang, Yang Liu, Aihua Zheng, and Pingping Zhang. Decoupled feature-based mixture of experts for multi-modal object re-identification. *arXiv preprint arXiv:2412.10650*, 2024. 5, 6, 8
- [39] Yingquan Wang, Pingping Zhang, Dong Wang, and Huchuan Lu. Other tokens matter: Exploring global and local features of vision transformers for object re-identification. *CVIU*, 244:104030, 2024. 1
- [40] Yuhao Wang, Pingping Zhang, Xuehu Liu, Zhengzheng Tu, and Huchuan Lu. Unity is strength: Unifying convolutional and transformer features for better person re-identification. *TITS*, 2025. 1
- [41] Zi Wang, Chenglong Li, Aihua Zheng, Ran He, and Jin Tang. Interact, embed, and enlarge: Boosting modality-specific representations for multi-modal person re-identification. In *AAAI*, pages 2633–2641, 2022. 2, 5
- [42] Zi Wang, Huaibo Huang, Aihua Zheng, and Ran He. Heterogeneous test-time training for multi-modal person re-identification. In *AAAI*, pages 5850–5858, 2024. 2, 5, 6
- [43] Di Wu, Zhihui Liu, Zihan Chen, Shenglong Gan, Kaiwen Tan, Qin Wan, and Yaonan Wang. Lrmm: Low rank multi-scale multi-modal fusion for person re-identification based on rgb-ni-ti. *ESWA*, 263:125716, 2025. 5, 6
- [44] Zhuofan Xia, Xuran Pan, Shiji Song, Li Erran Li, and Gao Huang. Vision transformer with deformable attention. In *CVPR*, pages 4794–4803, 2022. 5
- [45] Bin Yang, Jun Chen, and Mang Ye. Shallow-deep collaborative learning for unsupervised visible-infrared person re-identification. In *CVPR*, pages 16870–16879, 2024. 1
- [46] Haoli Yin, Jiayao Li, Eva Schiller, Luke McDermott, and Daniel Cummings. Graft: Gradual fusion transformer for multimodal re-identification. *arXiv preprint arXiv:2310.16856*, 2023. 6, 5
- [47] Chenyang Yu, Xuehu Liu, Yingquan Wang, Pingping Zhang, and Huchuan Lu. Tf-clip: Learning text-free clip for video-based person re-identification. In *AAAI*, pages 6764–6772, 2024. 1
- [48] Jiazuo Yu, Haomiao Xiong, Lu Zhang, Haiwen Diao, Yunzhi Zhuge, Lanqing Hong, Dong Wang, Huchuan Lu, You He, and Long Chen. Llms can evolve continually on modality for x-modal reasoning. *arXiv preprint arXiv:2410.20178*, 2024. 2
- [49] Jiazuo Yu, Yunzhi Zhuge, Lu Zhang, Ping Hu, Dong Wang, Huchuan Lu, and You He. Boosting continual learning of vision-language models via mixture-of-experts adapters. In *CVPR*, pages 23219–23230, 2024. 2
- [50] Zhi Yu, Zhiyong Huang, Mingyang Hou, Jiaming Pei, Yan Yan, Yushi Liu, and Daming Sun. Representation selective coupling via token sparsification for multi-spectral object re-identification. *TCSVT*, 2024. 5, 6
- [51] Zhi Yu, Zhiyong Huang, Mingyang Hou, Yan Yan, and Yushi Liu. Wtsf-reid: Depth-driven window-oriented token selection and fusion for multi-modality vehicle re-identification with knowledge consistency constraint. *ESWA*, page 126921, 2025. 5, 6
- [52] Yajing Zhai, Yawen Zeng, Da Cao, and Shaofei Lu. Tri-reid: Towards multi-modal person re-identification via descriptive fusion model. In *ICMR*, pages 63–71, 2022. 2
- [53] Guowen Zhang, Pingping Zhang, Jinqing Qi, and Huchuan Lu. Hat: Hierarchical aggregation transformers for person re-identification. In *ACM MM*, pages 516–525, 2021. 1
- [54] Pingping Zhang, Yuhao Wang, Yang Liu, Zhengzheng Tu, and Huchuan Lu. Magic tokens: Select diverse tokens for multi-modal object re-identification. In *CVPR*, pages 17117–17126, 2024. 1, 2, 3, 5, 6
- [55] Qingyang Zhang, Yake Wei, Zongbo Han, Huazhu Fu, Xi Peng, Cheng Deng, Qinghua Hu, Cai Xu, Jie Wen, Di Hu, et al. Multimodal fusion on low-quality data: A comprehensive survey. *arXiv preprint arXiv:2404.18947*, 2024. 1
- [56] Aihua Zheng, Zi Wang, Zihan Chen, Chenglong Li, and Jin Tang. Robust multi-modality person re-identification. In *AAAI*, pages 3529–3537, 2021. 2, 5, 6
- [57] Aihua Zheng, Ziling He, Zi Wang, Chenglong Li, and Jin Tang. Dynamic enhancement network for partial multi-modality person re-identification. *arXiv preprint arXiv:2305.15762*, 2023. 2, 5
- [58] Aihua Zheng, Xianpeng Zhu, Zhiqi Ma, Chenglong Li, Jin Tang, and Jixin Ma. Cross-directional consistency network with adaptive layer normalization for multi-spectral vehicle re-identification and a high-quality benchmark. *Information Fusion*, 100:101901, 2023. 2, 5, 6
- [59] Aihua Zheng, Zhiqi Ma, Yongqi Sun, Zi Wang, Chenglong Li, and Jin Tang. Flare-aware cross-modal enhancement network for multi-spectral vehicle re-identification. *Information Fusion*, 116:102800, 2025. 6
- [60] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random erasing data augmentation. In *AAAI*, pages 13001–13008, 2020. 6

IDEA: Inverted Text with Cooperative Deformable Aggregation for Multi-modal Object Re-Identification

Supplementary Material

6. Introduction

In this supplementary material, we provide comprehensive experimental details, visual examples and extended analyses to support the findings of the main manuscript. To be specific, the supplementary material is organized as follows:

1. **Detailed description of our caption generation pipeline and more examples from different datasets:**
 - Detailed description of the caption generation pipeline
 - More examples from the person/vehicle ReID datasets
2. **Details of the proposed modules and experiments:**
 - Implementation details of the proposed modules
 - More explanations of models in the ablation study
3. **Module validation and hyper-parameter analysis:**
 - Analysis of model parameters and performance
 - Comparison with different inverse directions in IMFE
 - Comparison with CLIP-based method
 - Training efficiency comparison
 - Effect of text with different modalities combination
 - Exploration of extreme cases with missing text
 - Effect of key modules on vehicle dataset
 - Hyper-parameter analysis for the vehicle dataset
4. **Visualization analysis of IDEA:**
 - Visualization of multi-modal ranking list
 - Visualization of the channel activation maps

These analyses provide a deeper understanding of our constructed datasets, the proposed modules and experimental results, further validating the effectiveness of our method.

7. Multi-modal Caption Generation

7.1. Details of the Caption Generation Pipeline

To bridge the gap in multi-modal object ReID captioning, we propose a novel caption generation pipeline that leverages MLLMs to generate informative text descriptions. Specifically, our pipeline consists of two steps: (1) Caption Generation and (2) Attribute Extraction. Without loss of generality, we take the multi-modal person ReID dataset as an example to illustrate the caption generation pipeline.

Caption Generation: The unique characteristics of multi-spectral images and the limitations of MLLMs in captioning make simultaneous annotation of paired multi-spectral images challenging. When presented with images from multiple spectra, MLLMs often focus on shared semantic information while neglecting modality-specific details. To address this, we design modality-specific templates to annotate each spectrum independently, ensuring rich and detailed semantic descriptions. Specifically, this caption

generation step offers two main advantages. First, it enhances modality-specific descriptions by simplifying the captioning task, enabling the model to focus on detailed attributes unique to each spectrum. Second, it broadens applicability by supporting more cases such as text-to-image ReID in the infrared spectrum, which is particularly advantageous in challenging visual environments. Technically, our modality-specific templates consist of three key components: a modality-specific prefix, a generic annotation template and an anti-hallucination instruction.

1. **Modality-specific Prefix:** Directs the model to focus on a particular spectrum (e.g., RGB, NIR, or TIR), providing a detailed prompt to describe key attributes such as clothing, hairstyle and belongings:

“Write a comprehensive description of the person’s overall appearance based on the [RGB/NIR/TIR] image, strictly following this template. Include the following attributes: ‘upper garment’, ‘lower garment’, ‘shoes’, ‘hairstyle’, ‘gender’, ‘age group’ and ‘belongings’. Use specific details, including color, patterns and texture details. Please follow this structure: ”

2. **Generic Annotation Template:** Ensures consistency across captions with the following structure:

“The {Gender} is wearing a {Upper.Color} {Upper.Garment} with {Lower.Color} {Lower.Garment} and {Shoe.Color} {Shoes}. The {Gender} has {Hairstyle} {Hair.Color} hair and appears to be {Age.Group}. The {Gender} is carrying {Belongings}.”

3. **Anti-hallucination Instruction:** Prevents imagined details [32] by explicitly guiding the model to focus only on visible attributes:

“If certain attributes are not visible, ignore them. Do not imagine contents not present in the image. Adhere strictly to the format without adding extra explanations.”

By combining images from specific spectra with their corresponding templates, MLLMs generate tailored and comprehensive text descriptions for each modality. This approach provides detailed semantic information and effectively addresses the challenges of multi-modal data annotation.

Attribute Extraction: During the Caption Generation phase, we generate informative text descriptions for each spectrum. However, due to the inherent randomness of MLLMs during the annotation process, issues such as disorganized structure and overly verbose descriptions persist,




Figure 10. More examples from the multi-modal person ReID dataset RGBNT201.

even with the use of strict templates during generation. To address these challenges, we initially use regular expressions to extract key information from the sentences. While this approach captures most essential attributes, it struggles

with handling diverse and complex sentence structures. Fortunately, MLLMs exhibit robust capabilities in extracting key information from text. Leveraging this, we feed the generated descriptions back into the MLLMs with an at-




The vehicle is a silver SUV with a Volvo logo visible on its front grille. Its license plate is clearly readable as "A82730".

There are no noticeable damages or modifications, but there might be a slight scratch on the rear bumper. No additional accessories like a bike rack or trailer hitch were observed. Headlights appear off.



The vehicle is a silver SUV with a Volvo logo visible on its front grille. Its license plate is partially obscured but appears to be white with black lettering. There are no noticeable damages or modifications to the car, and there are no additional accessories or customization seen. The headlights appear to be off, indicating nighttime or low-light conditions.



The vehicle is a dark-colored SUV with a Volvo logo visible on its front grille. Its license plate is obscured from view. There are no noticeable damages or modifications to the car, but there are two small dents on its rear bumper. Additionally, there seems to be a bike rack attached to the back of the SUV. The lighting status cannot be determined.



The vehicle in question is a gray sedan with a black radiator grille and white headlamps. Although its license plate number seems to read "AU4865" due to some leaves obstructing it, there's no visible damage or modification on the car. Additionally, there aren't any extra accessories that stand out, and the headlights remain turned off.



The vehicle is a gray car with a black grille and wheels. Its license plate is partially obscured by tree branches. There are no visible damages or modifications, but there may be additional accessories such as mirrors or spoilers that cannot be seen clearly from this angle. The headlights appear to be off.



The vehicle is white and appears to be an SUV. There is no clear view of the license plate due to its position behind the car. There are no visible damages or modifications to the exterior of the vehicle. However, there are two small stickers attached to the rear window. The headlights are off but the brake lights appear to be illuminated.



The vehicle is a blue SUV with a "T1257" license plate located at the rear. There are no visible damages or modifications, but there may be additional accessories such as a roof rack or bike holder. The lights appear to be off.



The vehicle is a gray SUV with a Mazda logo. Its license plate is partially obscured but reads "T1257". There are no other discernible features such as accessories or additional modifications. The lights appear to be off.



The vehicle is a white SUV with a Mazda logo. Its license plate is partially obscured but appears to be blue with white lettering. There are no visible damages or modifications. The headlights appear to be off, suggesting either early morning or late evening hours.

Figure 11. More examples from the multi-modal vehicle ReID dataset MSVR310.

tribute extraction prompt to specify the predefined attributes to be extracted. The extracted information is then mapped into the template we use in the caption generation phase:

"Extract the key attributes from the sentence I give you and fill them into the following template: The {Gender} is wearing a {Upper.Color} {Upper.Garment} with {Lower.Color} {Lower.Garment} and {Shoe.Color} {Shoes}. The {Gender} has {Hairstyle} {Hair.Color} hair and appears to be {Age.Group}. The {Gender} is carrying {Belongings}. Strictly follow the template, do not add any extra information."

This process ensures the generation of concise and informative textual descriptions, effectively mitigating the inconsistencies introduced during the initial phase.

7.2. More Examples from Different Datasets

Multi-modal Person ReID Dataset Examples. In Fig. 10, we present additional examples from the RGBNT201 dataset to demonstrate the effectiveness of our caption generation pipeline. The generated text descriptions provide rich and informative attributes for each spectrum, including clothing, hairstyle and belongings. These examples verify the effectiveness of our multi-modal captioning pipeline in

producing detailed descriptions, which are crucial for multi-modal person ReID under complex visual environments.

Multi-modal Vehicle ReID Dataset Examples. In Fig. 11, we present additional examples from the MSVR310 dataset to illustrate the effectiveness of our caption generation pipeline in the vehicle ReID scenario. The generated text descriptions capture detailed attributes for each spectrum, including vehicle type, color and license plate number. Notably, descriptions for RGB images reliably provide accurate license plate numbers, while NIR and TIR images focus on attributes such as vehicle type and logo. With these detailed and informative text descriptions, we can fully leverage the semantic information from text descriptions to enhance the performance of multi-modal vehicle ReID.

8. Details of Modules and Experiments

8.1. Details of the Proposed Modules in IDEA

Details of Modal Prefixes. To differentiate textual descriptions from various modalities while fine-tuning the CLIP text encoder, we introduce Modal Prefixes. These prefixes consist of two components: a fixed textual description highlighting the characteristics of a specific spectrum and learnable tokens for fine-tuning. This combination enables modality-aware embedding representations. To be specific, the Modal Prefixes for each spectrum are defined as follows:

- RGB: “An image of a XXXX person in the visible spectrum, capturing natural colors and fine details: ”
- NIR: “An image of a XXXX person in the near infrared spectrum, capturing contrasts and surface reflectance: ”
- TIR: “An image of a XXXX person in the thermal infrared spectrum, capturing heat emissions as temperature gradients: ”

After tokenizing the input text and converting them into embeddings, we add randomly initialized learnable tokens to the embedding at the position corresponding to “XXXX”. During fine-tuning, these learnable tokens are updated to capture semantic information among the text descriptions, providing more information for the subsequent InverseNet.

Details of InverseNet. To fully leverage the semantic information from text descriptions, we propose the InverseNet. Specifically, InverseNet $\mathcal{I}$ is a simple MLP layer that takes the text feature $\hat{f}_m^t$ as input and outputs a pseudo token f_m^t . The input text feature $\hat{f}_m^t$ is determined based on the presence of learnable tokens in the Modal Prefixes:

- If no learnable tokens are added, we directly use the global feature f^{end} from the “endoftext” index position of the text, following previous works:

$$\hat{f}_m^t = f^{\text{end}}. \quad (12)$$

- If learnable tokens are present, we extract the feature corresponding to the position of “XXXX”, denoted as $f^{\text{prompt}} \in \mathbb{R}^{N_p \times C}$ and average them with f^{end} as follows:

$$\hat{f}_m^t = \frac{1}{N_p + 1} \left(f^{\text{end}} + \sum_{i=1}^{N_p} f_i^{\text{prompt}} \right). \quad (13)$$

This approach ensures that the text representation incorporates both the global context from the “endoftext” index and the other semantic information provided by the learnable tokens. Then, the pseudo token f_m^t is calculated as follows:

$$f_m^t = \mathcal{I}(\hat{f}_m^t) = \omega(\varphi(\omega(\delta(\varphi(\hat{f}_m^t))))), \quad (14)$$

where φ , δ and ω denote the linear layer, GeLU activation function and dropout layer, respectively. Then, the pseudo token f_m^t is concatenated with the visual features, guiding the visual encoder to focus on the interaction between the semantic information and image details.

Details of CDA. To further enhance the interaction between global features and discriminative local information, we propose CDA. In this supplementary material, we clarify the following key points regarding CDA:

- **Handling Sampling Point Overflow.** During the process of feature extraction, predicted sampling points $\hat{P}$ may occasionally fall outside the bounds of the feature map. To ensure that all sampled locations remain within the valid range, the coordinates are first normalized to the range $[-1, +1]$. Subsequently, any sampling point that exceeds this range is clipped to the closest valid value. Mathematically, for the sampling point $(\hat{x}, \hat{y})$, the clipping process is defined with the following equations:

$$\hat{x} = \text{clip}(\hat{x}, -1, +1), \quad \hat{y} = \text{clip}(\hat{y}, -1, +1). \quad (15)$$

This clipping process ensures that all coordinates stay within the valid range and allows for stable sampling, particularly at the boundaries of the feature map.

- **Bilinear Interpolation for Feature Sampling.** Once the sampling points are determined, we employ bilinear interpolation to extract features from the original feature map $\hat{F}_m \in \mathbb{R}^{H \times W \times C}$. For each sampling point $(\hat{x}, \hat{y})$, we identify the four neighboring grid points surrounding it. Let (i, j) denote the top-left corner of the grid, and the neighboring grid points are then located at $(i + 1, j)$, $(i, j + 1)$, and $(i + 1, j + 1)$. The interpolation is based on the horizontal and vertical distances between the sampling point $\hat{P}$ and these neighboring grid points:

$$dx = \hat{x} - i, \quad dy = \hat{y} - j. \quad (16)$$

Then, we calculate the bilinear interpolation weights as:

$$w_{mn} = (1 - m \cdot dx)(1 - n \cdot dy), \quad m, n \in \{0, 1\}. \quad (17)$$

The feature value at the sampled point is then computed as a weighted sum of its neighboring grid points:

$$\begin{aligned} \hat{F}_m(\hat{x}, \hat{y}) = & w_{00}\hat{F}_m(i, j) + w_{10}\hat{F}_m(i + 1, j) + \\ & w_{01}\hat{F}_m(i, j + 1) + w_{11}\hat{F}_m(i + 1, j + 1). \end{aligned} \quad (18)$$

This bilinear interpolation procedure ensures that the sampled features are smooth and accurate, effectively preserving discriminative local information. Finally, we get the sampled features from different modalities and re-shape them into the token format, which is then concatenated to get $F_S \in \mathbb{R}^{3N_S \times C}$, where $N_S = H_S \times W_S$.

With the above details of our proposed modules, we can effectively leverage the semantic information from text descriptions and enhance the interaction between global features and discriminative local information, leading to superior performance in multi-modal object ReID tasks.

8.2. Details of Models in the Ablation Study

Model B in Tab. 3. In Tab. 3 of the main manuscript, Model B integrates text information using a parallel structure. Specifically, this parallel structure is depicted in Fig. 12 (a), where the text feature is directly concatenated with the visual feature for retrieval. In this configuration, the text information is not utilized as effectively as in IMFE.

Model A in Tab. 5. In Tab. 5 of the main manuscript, Model A does not any components in CDA. All the patch tokens from different modalities are directly concatenated to interact with each other using the self-attention mechanism. After the interaction, we pool the features from different modalities and concatenate them for retrieval. Due to the limited local information, the performance of Model A is inferior to the global feature in IMFE.

IDEA w/o Text in Tab. 6. To evaluate the importance of text information in IDEA, we compare the performance of IDEA with and without text information in Tab. 6. In fact, our IDEA can also work without text information. Under this setting, our IDEA is composed of **the visual baseline and the CDA**, which only utilizes visual information for retrieval. The results show that the text information significantly enhances the performance of IDEA, demonstrating the importance of leveraging semantic guidance.

IDEA w/o Offset in Tab. 7. In Tab. 7, we assess the effectiveness of the offset mechanism in CDA. When the offset is not used, we rely on convolutional layers to process $\hat{F}_m$, generating features that aggregate local information, which are of the same shape as our F_S . The results demonstrate that the offset mechanism enhances the capture of discriminative local information, leading to improved performance.

9. Module Validation and Analysis

Analysis of Model Parameters and Performance. In Tab. 8, we compare the trainable parameters and performance of our proposed IDEA model with several state-of-the-art methods, including CNN-based approaches and Transformer-based techniques. While Transformer-based methods typically have more parameters than CNN-based approaches, they often deliver superior performance due to




Figure 12. Comparison of different structures in the IMFE.

Methods	Params	RGBNT201		RGBNT100		MSVR310	
	M	mAP	R-1	mAP	R-1	mAP	R-1
HAMNet [18]	78.00	27.7	26.3	74.5	93.3	27.1	42.3
CCNet [58]	74.60	-	-	77.2	96.3	36.4	55.2
IEEE [41]	109.22	49.5	48.4	-	-	-	-
GAFNet [10]	130.00	-	-	74.4	93.4	-	-
UniCat* [4]	259.02	57.0	55.7	79.4	96.2	-	-
GraFT* [46]	101.00	-	-	76.6	94.3	-	-
TOP-ReID* [37]	324.53	72.3	76.6	81.2	96.4	35.9	44.6
EDITOR* [54]	118.55	66.5	68.3	82.1	96.4	39.0	49.3
RSCNet* [50]	124.10	68.2	72.5	82.3	<u>96.6</u>	39.5	49.6
WTSF-ReID* [51]	143.60	67.9	72.2	82.2	96.5	39.2	49.1
MambaPro† [36]	74.20	78.9	83.4	83.9	94.7	<u>47.0</u>	56.5
DeMo† [38]	98.79	79.0	<u>82.3</u>	<u>86.2</u>	97.6	49.2	59.8
IDEA†	91.67	80.2	82.1	87.2	96.5	<u>47.0</u>	62.4

Table 8. Parameter and performance comparison with state-of-the-art methods. The best and second results are in bold and underlined, respectively. The symbol † denotes CLIP-based methods, * indicates ViT-based methods and others are CNN-based methods.

their strong generalization capabilities. Among these, our proposed IDEA stands out by achieving state-of-the-art performance with significantly fewer parameters. For instance, IDEA requires only 91.67M parameters, far less than TOP-ReID’s 324.53M, yet it achieves substantial improvements in mAP and Rank-1 accuracy. On the RGBNT201 dataset, IDEA achieves an mAP of 80.2%, surpassing the previous best of 79.0% mAP by DeMo. Similarly, on the RGBNT100 dataset, IDEA achieves an mAP of 87.2%, which is 1% higher than DeMo. These results demonstrate the effectiveness and efficiency of our proposed IDEA.

Comparison with Different Inverse Directions. To further validate the effectiveness of our proposed IMFE, we compare the performance of IDEA using different inverse directions on the RGBNT201 dataset. As shown in Tab. 9, Model B in Fig. 12 (c) outperforms Model A in Fig. 12

Index	Inverse Direction	Metrics			
		mAP	Rank-1	Rank-5	Rank-10
A	Image to Text	72.0	73.4	83.4	89.2
B	Text to Image	77.2	81.1	88.4	92.2

Table 9. Comparison of inverse directions on RGBNT201.

Methods	Params	Metrics			
	M	mAP	Rank-1	Rank-5	Rank-10
TOP-ReID (CLIP)	324.53	<u>73.3</u>	<u>77.2</u>	<u>85.9</u>	<u>90.1</u>
IDEA	91.67	80.2	82.1	90.0	93.3

Table 10. Comparison with TOP-ReID (CLIP) on RGBNT201.

Model	Training Time (h)	Memory (GB)	Time per Epoch (min)	Samples/s
TOP-ReID	0.6650	17.80	0.3325	168.03
EDITOR	0.4074	16.41	0.3492	167.01
IDEA	0.3075	18.02	0.3600	159.68

Table 11. Training efficiency comparison on RGBNT201.

(b), achieving an mAP of 77.2% and Rank-1 accuracy of 81.1%, compared to 72.0% mAP and 73.4% Rank-1 accuracy for the image-to-text direction. These findings highlight the importance of leveraging semantic guidance from text descriptions to enrich visual interactions. However, the inherent limitations of multi-modal captioning, where the text modality serves as the primary fusion source, may introduce inconsistencies across spectra, resulting in a performance drop in certain cases. Thus, we adopt the text-to-image direction in our IDEA model as default setting.

Comparison with CLIP-based TOP-ReID. In Tab. 10, we compare our IDEA model with TOP-ReID [37], both utilizing the CLIP vision encoder. IDEA consistently outperforms TOP-ReID on the RGBNT201 dataset, achieving an mAP of 80.2% and Rank-1 accuracy of 82.1%, compared to TOP-ReID’s 73.3% mAP and 77.2% Rank-1 accuracy. These results highlight the effectiveness of our proposed modules in fully leveraging CLIP’s knowledge.

Training Efficiency Comparison. As shown in Tab. 11, IDEA demonstrates competitive training efficiency. Moreover, it requires **less time to converge**, benefiting from the effective utilization of semantic information from text descriptions. This guidance enables the model to focus on discriminative local features, accelerating convergence. Furthermore, GPU memory usage remains **below 24GB** during both training and inference. On the smaller RGBNT201 dataset, training can be completed **within 20 minutes**, further highlighting the efficiency of our proposed IDEA.

Effect of Text with Different Modalities Combination. In Tab. 12, we compare the performance of different modality combinations on the RGBNT201 dataset. Here, **R** refers to the RGB image modality along with text annotations derived from RGB images, while **N** and **T** represent the NIR and TIR image-text pairs, respectively. A crucial aspect of multi-modal object ReID is effectively integrating comple-

Modality	R	N	T	R+N	R+T	N+T	R+N+T
mAP	39.9	27.1	43.3	58.4	71.5	62.9	80.2

Table 12. Comparison of modalities combination on RGBNT201.

Cases	Full	M(R)	M(N)	M(T)	M(RN)	M(RT)	M(NT)	M(RNT)
mAP	80.2	79.8	80.0	79.9	79.7	79.6	79.8	79.5

Table 13. Comparison of text-missing cases on RGBNT201. M(X) indicates that the text annotations for modality X are blank strings.

Index	Modules			Metrics	
	Text	IMFE	CDA	mAP	Rank-1
A	×	×	×	40.4	56.0
B	✓	×	×	43.5	57.9
C	✓	✓	×	<u>45.7</u>	<u>61.1</u>
D	✓	✓	✓	47.0	62.4

Table 14. Comparison with different modules on MSVR310.

mentary information across modalities. Even with the addition of textual annotations, multi-spectral information remains indispensable. As shown in Tab. 12, incorporating IR data significantly enhances performance, underscoring its critical role in capturing essential identity cues. The results highlight that while text provides valuable semantic guidance, it cannot fully replace the rich visual and spectral details contributed by multi-modal fusion.

Exploration of Extreme Cases with Missing Text. In Tab. 13, we explore the impact of missing text annotations during retrieval on the RGBNT201 dataset. Notably, in the most extreme case where all text annotations are removed, IDEA still achieves a mAP of 79.5%, which is only a marginal drop from the fully annotated setting. This indicates that during training, the semantic guidance provided by text effectively enhances the fusion between multi-modal features, enabling the model to learn more robust representations. These findings validate that IDEA does not excessively rely on textual input during inference but instead leverages structured multi-modal learning to develop a more robust and adaptive feature representation.

Effect of Key Modules on Vehicle Dataset. In Tab. 14, we present the performance comparison of different modules on the MSVR310 vehicle dataset. Model A, which only uses visual information, achieves an mAP of 40.4% and Rank-1 accuracy of 56.0%. Model B introduces the text information with a parallel structure, leading to a 3.1% improvement in mAP compared to Model A. Model C further incorporates the IMFE, which fully leverages the semantic guidance from the text, resulting in a 2.2% improvement in mAP. Finally, Model D integrates the CDA, leading to a 1.3% improvement in mAP compared to Model C, achieving the best performance with an mAP of 47.0% and Rank-1 accuracy of 62.4%. These results fully validate the effectiveness and efficiency of our proposed modules.




Figure 13. Performance comparison with different prompt lengths and offset factors on multi-modal vehicle dataset MSVR310.

Hyper-parameter Analysis for Vehicle Dataset. Fig. 13 presents a performance comparison across different prompt lengths and offset factors on the MSVR310 vehicle dataset. The results demonstrate a general performance improvement with increasing prompt length. However, performance begins to decline when the prompt length exceeds 1. Consequently, we adopt a prompt length of 1 for subsequent experiments on MSVR310. A similar observation holds for the RGBNT201 dataset, where longer prompts truncate text descriptions, leading to the loss of critical information. For the offset factor, the performance remains relatively stable, with the optimal result achieved at an offset factor of 25.

10. Visualization Analysis of IDEA

10.1. More Visualization for Person ReID

Visualization of Channel Activation Maps on the Person ReID Dataset. Fig. 14 illustrates the channel activation maps across different modalities for person ReID datasets within our IDEA framework. Each modality exhibits distinct activation patterns, reflecting their unique spectral properties. Notably, these maps successfully identify semantic regions, including hair, clothing and accessories, underscoring the capability of our proposed modules to effectively leverage multi-modal information.

Visualization of Multi-modal Ranking List with Different Modules in IDEA. In Fig. 16, we present the multi-modal ranking list comparison with different modules on the RGBNT201 dataset. The baseline model without text information exhibits a significant number of incorrect matches, as shown in Fig. 16 (a). By incorporating text information, the model can better distinguish between correct and incorrect matches, as shown in Fig. 16 (b). The introduction of IMFE further enhances the model’s ability to identify correct matches, as shown in Fig. 16 (c). Finally, the integration of CDA significantly improves the model’s performance, as shown in Fig. 16 (d). These results demonstrate the effectiveness of our proposed modules in enhancing the performance of multi-modal object ReID.

Visualization of Multi-modal Ranking List Comparison




Figure 14. Visualization of channel activation maps of different modalities on the person ReID dataset.




Figure 15. Visualization of channel activation maps of different modalities on the vehicle ReID dataset.

with TOP-ReID. In Fig. 17, we present the multi-modal ranking list comparison with TOP-ReID on the RGBNT201 person dataset. We choose hard examples that TOP-ReID struggles with and compare the performance of our IDEA model. The results verify that our IDEA model consistently outperforms TOP-ReID, achieving accurate matches.

10.2. More Visualization for Vehicle ReID

Visualization of Channel Activation Maps on the Vehicle ReID Dataset. In Fig. 15, we present the channel activation maps of different modalities in IDEA on the vehicle ReID dataset. Different modalities focus on distinct semantic regions, such as the corners of the vehicle and the license plate. These activation maps effectively capture discriminative local information, highlighting the importance of leveraging interaction between global features and discriminative local information in multi-modal vehicle ReID.

Visualization of Multi-modal Ranking List with Different Modules. Fig. 18 provides a comparative analysis of multi-modal ranking lists across different module configurations on the MSVR310 vehicle dataset. The baseline model, as seen in Fig. 18 (a), relies solely on visual features and fails to reliably distinguish between vehicles with similar appearances, leading to frequent mismatches. Integrating textual information into the model, as shown in Fig. 18




Figure 16. Ranking list comparison with different modules on the person ReID dataset RGBNT201. (a) Baseline. (b) Baseline + Text. (c) Baseline + IMFE. (d) Baseline + IMFE + CDA. The green box indicates the correct match, while the red box indicates the incorrect match.



Figure 17. Multi-modal ranking list comparison with TOP-ReID on the person ReID dataset RGBNT201. (a) TOP-ReID. (b) IDEA. The multi-modal ranking list of TOP-ReID is sourced from the supplementary material of DeMo [38].




Figure 18. Ranking list comparison with different modules on the vehicle ReID dataset MSVR310. (a) Baseline. (b) Baseline + Text. (c) Baseline + IMFE. (d) Baseline + IMFE + CDA.




Figure 19. Multi-modal ranking list comparison with EDITOR on the vehicle ReID dataset MSVR310. (a) EDITOR. (b) IDEA.

(b), greatly enhances the retrieval process, enabling the system to utilize semantic cues for better discrimination. With the introduction of IMFE (Fig. 18 (c)), the model demonstrates an ability to capture more nuanced modality-specific information, resulting in a marked improvement in retrieval accuracy. Finally, the addition of CDA (Fig. 18 (d)) further refines the ranking results, achieving a level of performance that effectively balances semantic understanding and visual precision. These findings validate the effectiveness of our modules in addressing the challenges of vehicle ReID.

Visualization of Multi-modal Ranking List Comparison with EDITOR. In Fig. 19, we compare our IDEA model

against EDITOR, the previous state-of-the-art approach on the MSVR310 dataset. EDITOR, while highly effective in many scenarios, struggles with difficult cases involving visually similar vehicles or missing modality-specific details. This limitation is evident in its inability to consistently rank correct matches, especially when semantic guidance is needed to resolve ambiguities. In contrast, our IDEA model leverages advanced multi-modal modules to seamlessly integrate visual and textual cues, resulting in a dramatic reduction in errors even for the hard cases. This comparative analysis highlights IDEA’s superiority over EDITOR, verifying its robustness in tackling real-world challenges.