

Hierarchical Prompt Learning for Image- and Text-Based Person Re-Identification

Linhan Zhou^{1*}, Shuang Li^{2*}, Neng Dong³, Yonghang Tai⁴, Yafei Zhang¹, Huafeng Li^{1†}

¹Faculty of Information Engineering and Automation, Kunming University of Science and Technology

²School of Computer Science and Technology, Chongqing University of Post and Telecommunication

³School of Computer Science and Engineering, Nanjing University of Science and Technology

⁴School of Physics and Electronic Information, Yunnan Normal University

LinhanZhouUltra@outlook.com, shuangli936@gmail.com, neng.dong@njust.edu.cn, taiyonghang@126.com, {zyfeimail, hfchina99}@163.com

Abstract

Person re-identification (ReID) aims to retrieve target pedestrian images given either visual queries (image-to-image, I2I) or textual descriptions (text-to-image, T2I). Although both tasks share a common retrieval objective, they pose distinct challenges: I2I emphasizes discriminative identity learning, while T2I requires accurate cross-modal semantic alignment. Existing methods often treat these tasks separately, which may lead to representation entanglement and suboptimal performance. To address this, we propose a unified framework named Hierarchical Prompt Learning (HPL), which leverages task-aware prompt modeling to jointly optimize both tasks. Specifically, we first introduce a Task-Routed Transformer, which incorporates dual classification tokens into a shared visual encoder to route features for I2I and T2I branches respectively. On top of this, we develop a hierarchical prompt generation scheme that integrates identity-level learnable tokens with instance-level pseudo-text tokens. These pseudo-tokens are derived from image or text features via modality-specific inversion networks, injecting fine-grained, instance-specific semantics into the prompts. Furthermore, we propose a Cross-Modal Prompt Regularization strategy to enforce semantic alignment in the prompt token space, ensuring that pseudo-prompts preserve source-modality characteristics while enhancing cross-modal transferability. Extensive experiments on multiple ReID benchmarks validate the effectiveness of our method, achieving state-of-the-art performance on both I2I and T2I tasks.

Code — <https://github.com/LH-Z-Ac/HPL-AAAI26>

Introduction

Person Re-Identification (ReID) aims to retrieve a target individual from large-scale visual data given a query, and plays a key role in surveillance and public security. Based on query modality, ReID is divided into Image-to-Image (I2I) and Text-to-Image (T2I) (Li et al. 2017) tasks. I2I focuses on extracting identity-discriminative features robust to viewpoint

*Equal Contribution.

†Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

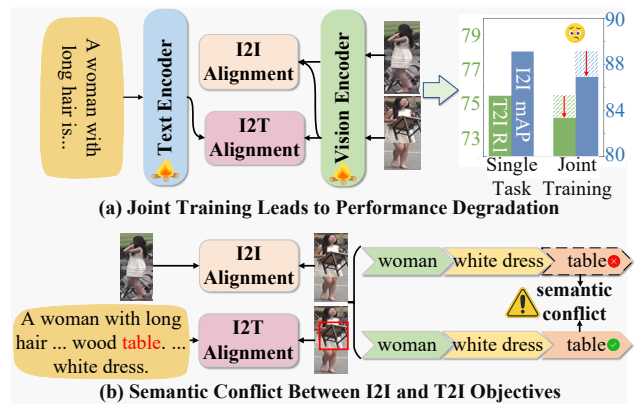


Figure 1: (a) Performance degradation occurs when jointly training I2I and T2I ReID tasks in a single model, compared to training each task independently. (b) The underlying cause lies in the semantic conflict: T2I emphasizes instance-specific attributes (e.g., “holding a table”) highlighted in text but ignored by I2I, despite shared identity-level cues such as clothing and gender.

and background changes, while T2I retrieves pedestrian images using natural language descriptions, enabling applications without image queries or requiring human interaction. Recent advancements have led to the development of CNN- and Vision-Transformer (Dosovitskiy et al. 2021; Yuan et al. 2025; Sun et al. 2024)-based models for I2I (Sun et al. 2018; Wang et al. 2018; Sun et al. 2020; He et al. 2021; Zheng et al. 2022), and fine-grained cross-modal alignment methods for T2I (Ma et al. 2022; Suo et al. 2022; Liu et al. 2025; Wang et al. 2021; Ji et al. 2022). However, most existing works treat these tasks in isolation, neglecting the heterogeneous nature of real-world applications where both image and text queries may co-exist. This motivates the urgent need for a unified framework that can simultaneously handle both I2I and T2I ReID tasks flexibly.

To achieve unified modeling for both I2I and T2I ReID, we attempt to train the two tasks within a single model jointly. However, as shown in Figure 1(a), this joint train-

ing strategy leads to a noticeable performance drop for both tasks compared to training them separately. We attribute this degradation to the semantic conflict between the two tasks. As illustrated in Figure 1(b), while I2I and T2I share identity-level semantics (e.g., clothing, gender), T2I additionally relies on instance-specific attributes (e.g., actions or objects) described in natural language. These attributes, such as the presence of a “table”, are often ignored by I2I alignment but emphasized in T2I supervision. This inconsistency introduces conflicting optimization signals that negatively affect both tasks, ultimately impairing the effectiveness of joint training.

Prompt learning has recently achieved remarkable success in image-to-image (I2I) person re-identification (ReID) by injecting explicit identity-level semantics into the representation learning process (Li, Sun, and Li 2023). While identity-level prompts are also beneficial for text-to-image (T2I) ReID, directly extending prompt learning to this task remains non-trivial. Unlike I2I, which primarily relies on stable identity cues, T2I additionally requires fine-grained, instance-specific semantics—such as actions, gestures, or contextual objects—that are described in natural language but often vary across samples and views. To bridge this gap, we draw inspiration from the prompt inversion mechanism proposed in GET (Wang et al. 2025b), which translates visual features into pseudo-textual prompts to facilitate cross-modal alignment. Building on this idea, we hypothesize that instance-level semantics can be directly inferred from either visual or textual inputs via modality-specific inversion networks. Based on this insight, we formulate a unified hierarchical prompt structure: “*A photo of [id-tokens] and [inst-tokens] person*”, where *[id-tokens]* represent learnable identity-level semantics, and *[inst-tokens]* encode dynamically generated instance-specific attributes. This formulation enables the model to jointly capture shared identity cues and task-specific semantic variations, thus accommodating the heterogeneous supervision signals of I2I and T2I ReID within a unified framework.

Motivated by this, we propose a unified ReID framework centered on Hierarchical Prompt Learning (HPL), which serves as the core mechanism to integrate identity-level stability with instance-level adaptability. HPL constructs structured prompt representations by combining the shared learnable identity prompt with dynamically generated instance prompt tokens, the latter derived from either visual or textual features via modality-specific inversion networks. These hierarchical prompts are injected into the CLIP backbone to provide task-aware semantic guidance for both I2I and T2I tasks. To support effective prompt-driven learning, we introduce two complementary modules. First, the Task-Routed Transformer (TRT) extends the CLIP visual encoder with dual classification tokens, enabling the model to process both tasks within a shared backbone while preserving task-specific objectives. Second, the Cross-Modal Prompt Regularization (CMPR) module aligns instance-level prompt tokens generated from different modalities in the textual prompt space, enhancing semantic consistency and improving cross-modal generalization. Together, these components form a cohesive architecture that enables unified modeling

across modalities and supervision types, effectively bridging the gap between I2I and T2I ReID within a single framework.

The main contributions of this paper are summarized as follows:

- We propose a unified person ReID framework that jointly handles image-to-image (I2I) and text-to-image (T2I) retrieval tasks within a single architecture.
- We design a task-aware prompting mechanism that integrates a Task-Routed Transformer (TRT) with a Hierarchical Prompt Learning (HPL) scheme. This allows the model to dynamically combine identity-level and instance-level semantics for both visual and textual queries.
- We introduce a Cross-Modal Prompt Regularization (CMPR) strategy to align instance-level prompt tokens across modalities, enhancing semantic consistency.
- Extensive experiments on standard ReID benchmarks demonstrate the effectiveness of our framework, achieving state-of-the-art performance on both I2I and T2I tasks.

Related Works

Image-to-Image Person Re-identification

Image-to-Image Person Re-identification (I2I ReID) retrieves pedestrian images across camera views using a query image. Early methods employed CNNs for feature encoding and similarity matching (Sun et al. 2018; Wang et al. 2018). Recently, Transformer-based models (Dosovitskiy et al. 2021; Sun et al. 2020; He et al. 2021; Zheng et al. 2022; Xia et al. 2023) have gained popularity due to their superior ability to capture global context.

More recently, the introduction of CLIP (Radford et al. 2021) has led to a surge in CLIP-based I2I ReID methods (Li, Sun, and Li 2023; Wang et al. 2025a; Liu et al. 2024a; Tong et al. 2025; Lin et al. 2023). While CLIP is inherently designed for cross-modal alignment, standard I2I ReID datasets lack textual annotations. To utilize CLIP’s full potential, these works construct pseudo-text descriptions for images to assist training. For example, (Li, Sun, and Li 2023) proposes a two-stage approach: prompt constructing and image-prompt alignment. Prompt-based methods have also been explored in text- (Yan et al. 2024; Li et al. 2024b) and clothes-changing-Re-ID (Wei et al. 2025). Building on this, (Wang et al. 2025a) introduces self-supervised strategies, such as masked text modeling and random image occlusion, to improve prompt quality. And (Liu et al. 2024a) further extracts body-shape-related phrases using CLIP.

Text-to-Image Person Re-identification

Text-to-Image Person Re-identification (T2I ReID) retrieves pedestrian images based on textual descriptions (Li et al. 2017). Early methods (Wang et al. 2020) decompose both image and text into attribute components and conduct a fine-grained matching. Recent approaches propose end-to-end frameworks, but still follow the earlier idea of fine-grained

alignment. Alignment strategies are typically explicit or implicit. Explicit methods, such as (Wang et al. 2021), perform alignment on global and local features separately. (Ji et al. 2022) addresses the asymmetry in cross-granularity alignment by introducing a non-symmetric alignment mechanism. Implicit methods, such as (Bai et al. 2023; Liu et al. 2025), guide the model to learn fine-grained alignments through masked language modeling and similar techniques.

Recently, CLIP-based methods have gained prominence in the T2I ReID domain (Wang et al. 2023; Jiang and Ye 2023; Wu et al. 2024; Tan et al. 2024). For example, (Wang et al. 2023) introduces a multi-completeness constraint and dynamic attribute prompts to preserve CLIP’s generalization. (Wu et al. 2024) incorporates a feature filtering mechanism to mitigate the effects of occlusion. Additionally, (Tan et al. 2024) pretrains a vision-language model for T2I ReID using large-scale unlabelled pedestrian images paired with synthetic textual descriptions generated by multi-modal language models. More recent studies have also explored leveraging large vision-language models for person re-identification (Niu et al. 2025).

However, most of these methods focus primarily on cross-modal alignment, while paying limited attention to intra-modal consistency. In contrast, our goal is to integrate I2I ReID with T2I ReID, thereby enabling an efficient multi-modal person retrieval system.

Our Method

We propose a unified framework named Hierarchical Prompt Learning, as shown in Figure 2. It comprises three modules: a Task-Routed Transformer (TRT) with dual classification tokens for task-specific encoding, a Hierarchical Prompt Learning (HPL) module for hierarchical prompt generation and alignment, and a Cross-Modal Prompt Regularization (CMPR) module to enforce modality consistency via instance-level prompt supervision.

Task-routed Transformer

Previous methods typically use separate ViT backbones for T2I and I2I person ReID, despite the shared identity-relevant semantics (e.g., clothing, body structure) between tasks. The class token in ViTs naturally aggregates contextual information via self-attention. This mechanism implies that semantics aggregated into the class token are guided by the task-dependent objectives. Motivated by this, we introduce a lightweight yet effective task-routed design by simply appending an additional class token into the CLIP visual encoder. Each class token is dedicated to one specific task: the original token v_{t2i}^i is optimized for T2I alignment, while the newly introduced token v_{i2i}^i focuses on identity discrimination in the I2I setting.

Specifically, given a visual input $V_i \in \mathbb{R}^{H \times W \times C}$ from either the T2I dataset \mathcal{D}_{t2i} or the I2I dataset \mathcal{D}_{i2i} , we feed it into the CLIP visual encoder \mathcal{M}^v , which is modified to include two class tokens. The resulting visual features are formulated as:

$$F_i^v = [v_{i2i}^i, v_{t2i}^i, v_1^i, \dots, v_N^i] = \mathcal{M}^v(V_i), \quad (1)$$

where v_{t2i}^i and v_{i2i}^i denote the classification tokens corresponding to the T2I and I2I tasks, respectively. The patch tokens v_1^i, \dots, v_N^i represent the spatial embeddings extracted from the image, where $N = H \times W / P^2$, and P denotes the patch size.

To enable the two classification tokens to specialize in their respective tasks, we introduce a multi-objective supervision strategy that guides each token through task-relevant losses:

$$\mathcal{L}_{base} = \mathcal{L}_{sdm} + \mathcal{L}_{id}^{t2i} + \mathcal{L}_{tri} + \mathcal{L}_{id}^{i2i} \quad (2)$$

Here, the T2I classification token v_{t2i}^i is optimized via the cross-modal similarity distribution matching loss \mathcal{L}_{sdm} and the cross-modal identity classification loss \mathcal{L}_{id}^{t2i} (Jiang and Ye 2023), both of which utilize the paired text features for cross-modal alignment. In parallel, the I2I classification token v_{i2i}^i is supervised by an identity classification loss \mathcal{L}_{id}^{i2i} and a triplet ranking loss \mathcal{L}_{tri} (Luo et al. 2019), promoting discriminative feature learning within the visual modality.

Hierarchical Prompt Learning

While TRT enables shared visual backbones, it lacks fine-grained semantic modeling. I2I benefits from identity-level prompts, as shown in CLIP-ReID (Li, Sun, and Li 2023), but T2I further relies on instance-specific details (e.g., clothing color, accessories). To address this, we propose Hierarchical Prompt Learning (HPL), which combines identity- and instance-level prompts to guide both tasks. HPL includes two modules: *Hierarchical Prompt Construction* and *Hierarchical Prompt Alignment*.

Hierarchical Prompt Construction. To provide unified textual supervision for both I2I and T2I person ReID, we design a hierarchical language template T_i : “A photo of [id-tokens] and [inst-tokens] person”, which differs from conventional prompt designs (e.g., CLIP-ReID). Here, the identity-level component [id-tokens] consists of a fixed number of learnable tokens, while the instance-level component [inst-tokens] is instantiated by pseudo-text tokens converted from either image or text features. This design injects both identity-level and instance-specific semantics into the prompt.

To enrich the prompt with fine-grained instance cues, we introduce two modality-specific inverse networks \mathcal{I}_v and \mathcal{I}_t to generate pseudo-text tokens from vision and text features, respectively. Given an image V_i and its associated prompt T_i , we compute the visual and textual features:

$$F_i^v = \mathcal{M}^v(V_i), \quad F_i^t = \mathcal{M}^t(T_i), \quad (3)$$

which are then fed into the corresponding inverse networks:

$$P_i^t = \mathcal{I}_t(F_i^t), \quad P_i^v = \mathcal{I}_v(F_i^v), \quad (4)$$

where \mathcal{I}_v and \mathcal{I}_t are composed of $N_{\mathcal{I}}$ Transformer layers following the ViT block architecture. These pseudo-text tokens are inserted into the hierarchical template to form: T_i^v : “A photo of [id-tokens] and P_i^v person.”; T_i^t : “A photo of [id-tokens] and P_i^t person.” The resulting hierarchical prompts are then encoded via another CLIP text encoder $\tilde{\mathcal{M}}^t$:

$$\tilde{v}_{eos}^i = \tilde{\mathcal{M}}^t(T_i^v), \quad \tilde{t}_{eos}^i = \tilde{\mathcal{M}}^t(T_i^t). \quad (5)$$

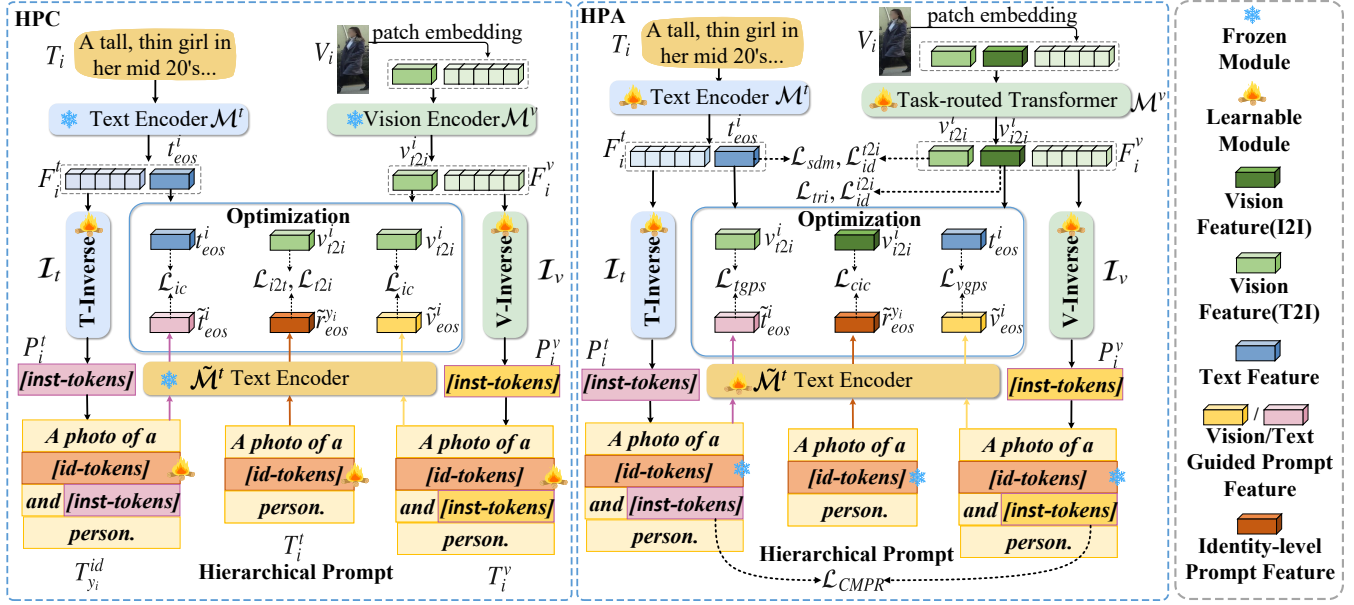


Figure 2: Overview of our framework. The framework comprises three core modules: (1) Task-Routed Transformer (TRT), which introduces dual classification tokens into the shared visual encoder to enable task-specific feature learning for I2I and T2I tasks; (2) Hierarchical Prompt Learning (HPL), which constructs and aligns identity-level and instance-level prompts via modality-specific inversion networks; (3) Cross-Modal Prompt Regularization (CMPR), which enforces semantic consistency between visual- and text-derived prompts at the instance level. Arrows indicate the directional flow of information and interactions among modules.

To ensure that these pseudo prompts faithfully retain modality-specific semantics, we apply an inversion consistency loss:

$$\mathcal{L}_{ic} = \frac{1}{|B|} \sum_{i \in B} \|\tilde{v}_{eos}^i - v_{t2i}^i\|_2^2 + \frac{1}{|B_{t2i}|} \sum_{i \in B_{t2i}} \|\tilde{t}_{eos}^i - t_{eos}^i\|_2^2, \quad (6)$$

where B_{t2i} and B indicate the batch for T2I and the entire batch, respectively, and $\|\cdot\|_2^2$ denotes the squared L2 norm. This supervision encourages the reconstructed prompts to reflect their input semantics while enhancing cross-modal transferability. During this stage, the encoders \mathcal{M}^v , \mathcal{M}^t and $\tilde{\mathcal{M}}^t$ are frozen, and only the inverse networks and prompts are updated.

To encode identity-level semantics, we instantiate a simplified version of the hierarchical prompt: $T_{y_i}^{id}$: “A photo of [id-tokens] person”, where [id-tokens] corresponds to the shared learnable identity-level tokens. This template is also passed the text encoder $\tilde{\mathcal{M}}^t$ to obtain the reference text representation:

$$\tilde{r}_{eos}^{y_i} = \tilde{\mathcal{M}}^t(T_{y_i}^{id}). \quad (7)$$

To guide identity-level alignment, we adopt two contrastive losses:

$$\mathcal{L}_{t2i} = -\frac{1}{|B|} \sum_{i \in B} \log \frac{\exp[\text{sim}(v_{t2i}^i, \tilde{r}_{eos}^{y_i})]}{\sum_{j \in B} \exp[\text{sim}(v_{t2i}^j, \tilde{r}_{eos}^{y_i})]}, \quad (8)$$

$$\mathcal{L}_{i2t} = -\frac{1}{|B|} \sum_{i \in B} \log \frac{\exp[\text{sim}(v_{i2i}^i, \tilde{r}_{eos}^{y_i})]}{\sum_{j \in B} \exp[\text{sim}(v_{i2i}^j, \tilde{r}_{eos}^{y_i})]}, \quad (9)$$

where $\text{sim}(\cdot, \cdot)$ denotes cosine similarity. These losses encourage the identity-level prompt to serve as a reliable semantic anchor during downstream training.

Hierarchical Prompt Alignment. In Hierarchical Prompt Alignment, we aim to utilize hierarchical prompts from Hierarchical Prompt Construction to help the model learn the necessary semantic features for T2I and I2I tasks. For the T2I task, we utilize the full hierarchical prompt—comprising both identity-level and instance-level components—to guide the visual and textual representations, respectively. Specifically, the visual classification token v_{t2i}^i is encouraged to align with the textual embedding \tilde{t}_{eos}^i derived from the pseudo-prompt $T_{y_i}^v$, while the textual feature t_{eos}^i is aligned with \tilde{v}_{eos}^i generated from the pseudo-prompt $T_{y_i}^t$. This is enforced through the instance-level prompt alignment loss \mathcal{L}_{ILPA} , which consists of text- and vision-guided prompts supervised loss \mathcal{L}_{tgps} and \mathcal{L}_{vgps} :

$$\begin{aligned} \mathcal{L}_{tgps} &= \frac{1}{|B_{t2i}|} \sum_{i \in B_{t2i}} \|\tilde{t}_{eos}^i - v_{t2i}^i\|_2^2 \\ \mathcal{L}_{vgps} &= \frac{1}{|B_{t2i}|} \sum_{i \in B_{t2i}} \|\tilde{v}_{eos}^i - t_{eos}^i\|_2^2 \\ \mathcal{L}_{ILPA} &= \mathcal{L}_{tgps} + \mathcal{L}_{vgps} \end{aligned} \quad (10)$$

For the I2I task, given a pedestrian image, we use the identity-level prompt $T_{y_i}^{id}$, which is encoded by the text encoder $\tilde{\mathcal{M}}^t$ to produce the textual representation $\tilde{r}_{eos}^{y_i}$. We then apply a cross-modal identity classification loss \mathcal{L}_{cic} to

supervise the alignment between the image feature v_{i2i}^i and the identity prompt:

$$\mathcal{L}_{cic} = - \sum_{i \in B} \log \frac{\exp[\text{sim}(v_{i2i}^i, \tilde{r}_{eos}^y)]}{\sum_{j=1}^{N_{id}} \exp[\text{sim}(v_{i2i}^i, \tilde{r}_{eos}^j)]} \quad (11)$$

where N_{id} denotes the total number of identities in the training set.

Advantages of HPL. Compared to previous prompt learning strategies that typically focus only on identity-level templates or rely on fixed textual embeddings, our hierarchical prompt scheme brings three key advantages: (1) *Fine-grained Adaptability*: The instance-level pseudo tokens capture unique instance-specific attributes (e.g., outfit, accessories), enabling the model to attend to nuanced visual cues beyond class identity. (2) *Bidirectional Cross-Modal Alignment*: By decoding visual/textual features into pseudo-text and enforcing both text-to-image and image-to-text alignment, our approach bridges modality discrepancies and improves semantic coherence across modalities. (3) *Prompt-based Harmonization*: Since the instance-level and identity-level prompts are concatenated together, the ILPA loss not only enhances the alignment of instance-specific semantics across modalities but also preserves identity consistency. These strengths allow the proposed HPL framework to generalize better across diverse ReID tasks, particularly under the challenging T2I scenario where textual guidance plays a crucial role.

Cross-Modal Prompt Regularization

While HPL captures both identity- and instance-level semantics, the alignment between visual and textual prompts remains implicit. The instance-level prompts P_i^v and P_i^t may encode modality-specific biases, causing semantic inconsistency—especially detrimental for T2I grounding. To mitigate this, we introduce cross-modal prompt regularization, which aligns pseudo-text tokens in the textual token space via inverse networks \mathcal{I}_v and \mathcal{I}_t directly. Formally, given an image-text pair (V_i, T_i) , we extract the corresponding instance-level prompts as $P_i^v = \mathcal{I}_v(F_i^v)$ and $P_i^t = \mathcal{I}_t(F_i^t)$, where $F_i^v = \mathcal{M}^v(V_i)$ and $F_i^t = \mathcal{M}^t(T_i)$ denote the visual and textual features, respectively. Then, the prompt-level alignment loss is defined as:

$$\mathcal{L}_{CMPR} = \frac{1}{|B_{t2i}|} \sum_{i \in B_{t2i}} \|P_i^t - P_i^v\|_F^2, \quad (12)$$

where $\|\cdot\|_F$ denotes the Frobenius norm, and B_{t2i} represents the batch of T2I training samples.

This regularization promotes a unified prompt representation for both modalities, ensuring image-guided prompts convey meaningful semantics that align with text-guided prompts, reducing cross-modal semantic drift and improving text-guided visual retrieval.

Optimization

The training process consists of two stages: constructing hierarchical prompts and leveraging them for representation learning.

Stage I: Prompt Construction. In this stage, we focus on learning discriminative and semantically consistent prompts. The contrastive losses \mathcal{L}_{t2i} and \mathcal{L}_{i2t} are applied to guide identity-level prompt learning, while the inversion consistency loss \mathcal{L}_{ic} ensures that the instance-level pseudo-prompts remain aligned with the original visual and textual features. The overall objective in this stage is:

$$\mathcal{L}_{construct} = \mathcal{L}_{t2i} + \mathcal{L}_{i2t} + \mathcal{L}_{ic}. \quad (13)$$

Stage II: Representation Learning. With the constructed prompts, we proceed to optimize cross-modal feature learning. The base loss \mathcal{L}_{base} supervises the I2I and T2I tasks. To enhance semantic alignment, we introduce identity-level and instance-level prompt alignment losses, \mathcal{L}_{cic} and \mathcal{L}_{ILPA} . Additionally, we incorporate the cross-modal prompt regularization loss \mathcal{L}_{CMPR} to further encourage consistency between visual- and text-derived prompts. The total loss for this stage is:

$$\mathcal{L}_{total} = \mathcal{L}_{base} + \mathcal{L}_{cic} + \lambda_1 \mathcal{L}_{ILPA} + \lambda_2 \mathcal{L}_{CMPR}. \quad (14)$$

This two-stage optimization facilitates joint learning of multi-level prompt representations and cross-modal features, effectively bridging the modality gap while preserving fine-grained identity semantics.

Experiments

Datasets and Evaluation Protocols

We evaluate our method on both image-to-image (I2I) and text-to-image (T2I) person re-identification tasks using three dataset combinations: *CUHK-PEDES + Market1501*, *ICFG-PEDES + MSMT17*, *RSTPReid + DukeMTMC-ReID*. *CUHK-PEDES* (Li et al. 2017) is a large-scale text-to-image benchmark, containing 40,208 images of 13,003 identities, each image annotated with two textual descriptions. *ICFG-PEDES* (Ding et al. 2021) contains 54,522 images from 4,102 identities, with each image paired with one natural language sentence. *RSTPReid* (Zhu et al. 2021) consists of 20,505 images from 4,101 identities captured by 15 surveillance cameras. *Market1501* (Zheng et al. 2015) is a widely used ReID benchmark, providing 32,668 images of 1,501 identities collected from 6 cameras. *MSMT17* (Wei et al. 2018) contains 126,441 images of 4,101 identities under various lighting conditions and viewpoints from 15 cameras. *DukeMTMC-ReID* (Ristani et al. 2016) consists of 36,411 images with 1,404 identities captured across 8 cameras. Following common protocols, we report Rank-1 accuracy and mean Average Precision (mAP) for all tasks.

Implementations

We adopt CLIP pretrained on LUPerson and large-scale synthetic image-text pairs as the backbone. Input images are resized to 384×128 , and textual inputs are truncated to 77 tokens. During the HPC stage, we train the model for 10 epochs using the Adam optimizer. The learning rates for the decoder and identity-level prompts are set to 5×10^{-5} and 0.02, respectively, with an exponential decay factor of 0.8 per epoch. Both the visual and textual inverse networks are configured with $N_{\mathcal{I}} = 4$ transformer layers. In the HPA

Method	Reference	CUHK-PEDES		ICFG-PEDES		RSTPReID	
		Rank-1	mAP	Rank-1	mAP	Rank-1	mAP
CFine (Yan et al. 2023)	TIP'23	69.57	-	60.83	-	50.55	-
IRRA (Jiang and Ye 2023)	CVPR'23	73.38	66.13	63.46	38.06	60.20	47.17
MDRL (Yang et al. 2024)	AAAI'24	74.56	-	65.88	-	-	-
TBPS-CLIP (Cao et al. 2024)	AAAI'24	73.54	65.38	65.05	39.83	61.95	48.26
IRLT (Liu et al. 2024b)	AAAI'24	74.46	-	64.72	-	61.49	-
UMSA (Zhao et al. 2024)	AAAI'24	74.25	66.15	65.62	38.78	63.40	49.28
LSPM (Li et al. 2024a)	TMM'24	74.38	67.74	64.40	42.60	-	-
FSRL (Wang et al. 2024)	ICMR'24	74.86	67.57	64.93	40.67	60.65	48.18
Propot (Yan et al. 2024)	MM'24	74.89	67.12	65.12	42.93	61.87	47.82
MMRef (Ma et al. 2025)	TMM'25	72.25	-	63.50	-	56.20	-
HPL(ours)	This Paper	76.28	70.90	66.61	44.14	64.00	53.13

Table 1: Comparison with T2I ReID methods on CUHK-PEDES, ICFG-PEDES and RSTPReID datasets. Experiments are conducted on combinations of T2I and I2I datasets. Specifically, the three T2I datasets in this table are respectively paired with the three I2I datasets used in Table 2.

Method	Reference	Market1501		MSMT17		DukeMTMC	
		Rank-1	mAP	Rank-1	mAP	Rank-1	mAP
CDNet (Li, Wu, and Zheng 2021)	CVPR'21	95.10	86.00	78.90	54.70	88.60	76.80
TransReID (He et al. 2021)	ICCV'21	95.20	88.90	85.30	67.40	90.70	82.00
DRL-Net (Jia et al. 2022)	TMM'22	94.70	86.90	78.40	55.30	88.10	76.60
DCAL (Zhu et al. 2022)	CVPR'22	94.70	87.50	83.10	64.00	89.00	80.10
CLIP-ReID (Li, Sun, and Li 2023)	AAAI'23	95.50	89.60	88.70	73.40	90.00	82.50
CLIP3DReID (Liu et al. 2024a)	CVPR'24	95.60	88.40	81.50	61.20	-	-
HPL(ours)	This Paper	95.99	89.82	91.04	79.01	90.35	82.93

Table 2: Comparison with I2I ReID methods on Market1501, MSMT17 and DukeMTMC datasets. The three I2I datasets in this table are respectively paired with the three T2I datasets used in Table 1.

Components			T2I		I2I	
TRT	HPL	CMPR	Rank-1	mAP	Rank-1	mAP
\times	\times	\times	74.22	70.45	94.50	86.91
\checkmark	\times	\times	75.27	70.80	95.36	88.98
\checkmark	\checkmark	\times	75.60	70.88	95.57	89.72
\checkmark	\checkmark	\checkmark	76.28	70.89	95.99	89.82

Table 3: Impact of TRT, HPL, and CMPR modules on Rank-1 and mAP on the CUHK-PEDES + Market1501 dataset.

stage, training proceeds for 60 epochs. A linear warm-up is applied in the first 5 epochs (from 10^{-6} to 10^{-5}), followed by cosine annealing. Each batch contains 64 image-text pairs from the T2I dataset and 64 images from the I2I dataset, with 4 instances per identity. To avoid data leakage and identity confusion, we remove test samples from training splits and unify identity labels across datasets. This ensures that samples belonging to the same person in different datasets are treated consistently during training. The loss weights λ_1 and λ_2 are set to 0.4 and 0.06, respectively. All experiments are conducted on a single NVIDIA RTX 4090 GPU.

Comparison with State-of-the-Art Methods

We report the performance of our unified framework on both T2I and I2I person Re-ID tasks across six benchmark

ILPA		T2I		I2I	
\mathcal{L}_{tgps}	\mathcal{L}_{vgps}	Rank-1	mAP	Rank-1	mAP
\times	\times	75.27	70.80	95.36	88.98
\checkmark	\times	75.58	70.84	95.42	89.35
\times	\checkmark	75.60	70.82	95.69	89.59
\checkmark	\checkmark	75.60	70.88	95.57	89.72

Table 4: Ablation of instance-level alignment using vision-guided, text-guided, and dual-modality prompt supervision on CUHK-PEDES + Market1501.

datasets.

T2I ReID. We evaluated our method on several T2I ReID datasets. As shown in Table 1, our method achieves competitive performance across all popular datasets. Specifically, our method obtains Rank-1 accuracy of 76.28%, 66.61%, and 64.00% on the CUHK-PEDES, ICFG-PEDES, and RSTPReID datasets, respectively. Moreover, unlike most existing methods, our approach simultaneously achieves strong performance on the I2I ReID task, highlighting its superiority in multi-task person retrieval.

I2I ReID. We also evaluated our method on several widely used I2I ReID benchmarks. As shown in Table 2, our approach outperforms CLIPReID (Li, Sun, and Li 2023), a




Figure 3: Effect of λ_1 and λ_2 on Rank-1 accuracy on CUHK-PEDES + Market1501.

state-of-the-art method that also leverages learnable prompts for training. Specifically, our method achieves Rank-1 accuracies of 95.99%, 91.04%, and 90.35% on Market1501, MSMT17, and DukeMTMC-ReID datasets, respectively. The corresponding mAP scores reach 89.82%, 79.01%, and 82.93%.

Ablation Studies and Analysis

To verify the effectiveness of each component of our method, we conducted ablation experiments on the CUHK-PEDES + Market1501 dataset combination. The results demonstrate that our method successfully mitigates the conflict between the two tasks, enabling better integration of I2I and T2I ReID.

Effectiveness of TRT. To evaluate the effectiveness of the Task-routed Transformer, we compared it with a baseline that uses a single class token. As shown in Table 3, introducing a dual class token design leads to a 1.07% improvement in T2I Rank-1 accuracy and a 2.07% increase in I2I mAP. These findings also suggest that T2I ReID and I2I ReID rely on different types of information. The dual-token structure helps decouple task-specific semantics.

Effectiveness of HPL. To evaluate the effectiveness of our Hierarchical Prompt Learning (HPL) module, we conduct ablation studies based on TRT with and without HPL. As shown in Table 3, introducing HPL yields considerable performance gains for both I2I and T2I tasks. This indicates that identity-level and sample-level prompts help the model capture more refined and task-aware features. However, the improvement remains limited when HPL is used alone.

We further investigate using uni-modal sample-level prompts from either the visual or textual modality. As shown in Table 4, both achieve gains over the baseline, confirming the benefit of sample-level guidance. Compared to uni-modal prompts, the dual-modal setting brings a slight drop in I2I Rank-1 (-0.12%) but improves mAP (+0.13%), suggesting more comprehensive identity coverage. The impact on T2I is marginal, likely due to sufficient grounding from the paired text.

Effectiveness of CMPR. Based on TRT and HPL, we conduct ablation studies with and without Cross-Modal Prompt Regularization (CMPR). As shown in Table 3, adding CMPR together with HPL leads to significant performance improvements across both I2I and T2I tasks (+1.01% on T2I Rank1 and +0.84% on I2I mAP). This demonstrates

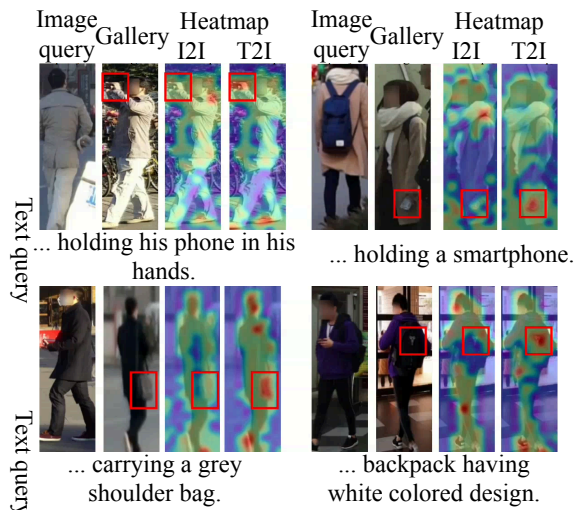


Figure 4: Grad-CAM comparison of I2I and T2I tokens, showing stronger focus on view-specific cues in T2I aligned with text descriptions.

that CMPR serves as an effective regularizer.

Hyperparameter Analysis. To investigate the effect of different values of λ_1 and λ_2 on model performance, we conducted a sensitivity analysis on the CUHK-PEDES + Market1501 dataset combination. As shown in Figure 3, both excessively small and large values of λ_1 and λ_2 lead to performance degradation. Based on experimental results, we set $\lambda_1 = 0.4$ and $\lambda_2 = 0.06$ to achieve optimal results.

Visualization Results. We visualize attention maps using Grad-CAM (Selvaraju et al. 2017) to analyze the behavior of I2I and T2I classification tokens (Figure 4). I2I tokens mainly focus on identity-related regions such as clothing and body shape, which remain consistent across views. In contrast, T2I tokens emphasize instance-specific details mentioned in text, e.g., phones or shoulder bags. For instance, in the second row, T2I attention highlights the grey shoulder bag while I2I ignores it. This modality-specific focus shows the model’s adaptive disentanglement of task-relevant semantics, alleviating supervision conflicts in joint training.

Conclusion

In this paper, we present a unified framework for image- and text-based person re-identification, jointly optimizing both tasks via task-aware prompt learning. A Task-Routed Transformer with dual classification tokens enables task-specific representation within a shared encoder. To capture multi-level semantics, a Hierarchical Prompt Learning module disentangles identity- and instance-level cues through modality-specific inversion. Furthermore, Cross-Modal Prompt Regularization aligns pseudo-prompts across modalities to reduce semantic inconsistency. Experiments on six benchmarks show that our method consistently surpasses state-of-the-art approaches on both I2I and T2I tasks.

Acknowledgments

This work was supported in part by the National Science Foundation of China under Grant 62276120 and Grant 61966021, the Yunnan Fundamental Research Projects under Grant 202301AV070004 and Grant 202401AS070106, the Major Science and Technology Special Projects of Yunnan Province under Grant 202502AD080006.

References

- Bai, Y.; Cao, M.; Gao, D.; Cao, Z.; Chen, C.; Fan, Z.; Nie, L.; and Zhang, M. 2023. RaSa: relation and sensitivity aware representation learning for text-based person search. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, 555–563.
- Cao, M.; Bai, Y.; Zeng, Z.; Ye, M.; and Zhang, M. 2024. An empirical study of clip for text-based person search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 465–473.
- Ding, Z.; Ding, C.; Shao, Z.; and Tao, D. 2021. Semantically Self-Aligned Network for Text-to-Image Part-aware Person Re-identification. [arXiv:2107.12666](https://arxiv.org/abs/2107.12666).
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.
- He, S.; Luo, H.; Wang, P.; Wang, F.; Li, H.; and Jiang, W. 2021. Transreid: Transformer-based object re-identification. In *Proceedings of the IEEE/CVF international conference on computer vision*, 15013–15022.
- Ji, Z.; Hu, J.; Liu, D.; Wu, L. Y.; and Zhao, Y. 2022. Asymmetric cross-scale alignment for text-based person search. *IEEE Transactions on Multimedia*, 25: 7699–7709.
- Jia, M.; Cheng, X.; Lu, S.; and Zhang, J. 2022. Learning disentangled representation implicitly via transformer for occluded person re-identification. *IEEE Transactions on Multimedia*, 25: 1294–1305.
- Jiang, D.; and Ye, M. 2023. Cross-modal implicit relation reasoning and aligning for text-to-image person retrieval. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2787–2797.
- Li, H.; Wu, G.; and Zheng, W.-S. 2021. Combined depth space based architecture search for person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6729–6738.
- Li, J.; Jiang, M.; Kong, J.; Tao, X.; and Luo, X. 2024a. Learning semantic polymorphic mapping for text-based person retrieval. *IEEE Transactions on Multimedia*.
- Li, S.; Sun, L.; and Li, Q. 2023. Clip-reid: exploiting vision-language model for image re-identification without concrete text labels. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 1405–1413.
- Li, S.; Xiao, T.; Li, H.; Zhou, B.; Yue, D.; and Wang, X. 2017. Person search with natural language description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1970–1979.
- Li, W.; Tan, L.; Dai, P.; and Zhang, Y. 2024b. Prompt decoupling for text-to-image person re-identification. *arXiv preprint arXiv:2401.02173*.
- Lin, Y.; Liu, C.; Chen, Y.; Hu, J.; Yin, B.; Yin, B.; and Wang, Z. 2023. Exploring part-informed visual-language learning for person re-identification. *arXiv preprint arXiv:2308.02738*.
- Liu, F.; Kim, M.; Ren, Z.; and Liu, X. 2024a. Distilling clip with dual guidance for learning discriminative human body shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 256–266.
- Liu, J.; Wei, D.; Liu, Y.; Zhang, S.; Yang, T.; Zhou, W.; Ding, W.; and Leung, V. C. M. 2025. SCMM: Calibrating Cross-modal Representations for Text-Based Person Search. [arXiv:2304.02278](https://arxiv.org/abs/2304.02278).
- Liu, Y.; Qin, G.; Chen, H.; Cheng, Z.; and Yang, X. 2024b. Causality-inspired invariant representation learning for text-based person retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 14052–14060.
- Luo, H.; Gu, Y.; Liao, X.; Lai, S.; and Jiang, W. 2019. Bag of tricks and a strong baseline for deep person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, 0–0.
- Ma, Y.; Xu, G.; Sun, X.; Yan, M.; Zhang, J.; and Ji, R. 2022. X-clip: End-to-end multi-grained contrastive learning for video-text retrieval. In *Proceedings of the 30th ACM international conference on multimedia*, 638–647.
- Ma, Z.; Chen, H.; Zeng, W.; Su, L.; and Zhang, S. 2025. Multi-modal Reference Learning for Fine-grained Text-to-Image Retrieval. *IEEE Transactions on Multimedia*, 1–14.
- Niu, K.; Yu, H.; Zhao, M.; Fu, T.; Yi, S.; Lu, W.; Li, B.; Qian, X.; and Xue, X. 2025. Chatreid: Open-ended interactive person retrieval via hierarchical progressive tuning for vision language models. *arXiv preprint arXiv:2502.19958*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Ristani, E.; Solera, F.; Zou, R.; Cucchiara, R.; and Tomasi, C. 2016. Performance measures and a data set for multi-target, multi-camera tracking. In *European conference on computer vision*, 17–35. Springer.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, 618–626.
- Sun, Y.; Cheng, C.; Zhang, Y.; Zhang, C.; Zheng, L.; Wang, Z.; and Wei, Y. 2020. Circle loss: A unified perspective of pair similarity optimization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6398–6407.

- Sun, Y.; Qin, Y.; Li, Y.; Peng, D.; Peng, X.; and Hu, P. 2024. Robust multi-view clustering with noisy correspondence. *IEEE Transactions on Knowledge and Data Engineering*.
- Sun, Y.; Zheng, L.; Yang, Y.; Tian, Q.; and Wang, S. 2018. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In *Proceedings of the European conference on computer vision (ECCV)*, 480–496.
- Suo, W.; Sun, M.; Niu, K.; Gao, Y.; Wang, P.; Zhang, Y.; and Wu, Q. 2022. A simple and robust correlation filtering method for text-based person search. In *European conference on computer vision*, 726–742. Springer.
- Tan, W.; Ding, C.; Jiang, J.; Wang, F.; Zhan, Y.; and Tao, D. 2024. Harnessing the power of mlms for transferable text-to-image person reid. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17127–17137.
- Tong, H.; Liu, J.; Wu, Y.; Zhao, G.; Zhang, F.; and Zha, Z.-J. 2025. Multi-granularity and Multi-modal Prompt Learning for Person Re-Identification. In *International Conference on Computational Visual Media*, 177–200. Springer.
- Wang, B.; Liang, Y.; Cai, L.; Huang, H.; and Zeng, H. 2025a. Image re-identification: Where self-supervision meets vision-language learning. *Image and Vision Computing*, 154: 105415.
- Wang, C.; Luo, Z.; Lin, Y.; and Li, S. 2021. Text-based person search via multi-granularity embedding learning. In *IJCAI*, 1068–1074.
- Wang, D.; Yan, F.; Wang, Y.; Zhao, L.; Liang, X.; Zhong, H.; and Zhang, R. 2024. Fine-grained Semantics-aware Representation Learning for Text-based Person Retrieval. In *Proceedings of the 2024 International Conference on Multimedia Retrieval*, 92–100.
- Wang, E.; Peng, Z.; Xie, Z.; Yang, F.; Liu, X.; and Cheng, M.-M. 2025b. Get: Unlocking the multi-modal potential of clip for generalized category discovery. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 20296–20306.
- Wang, G.; Yu, F.; Li, J.; Jia, Q.; and Ding, S. 2023. Exploiting the Textual Potential from Vision-Language Pre-training for Text-based Person Search. *arXiv e-prints*, arXiv–2303.
- Wang, G.; Yuan, Y.; Chen, X.; Li, J.; and Zhou, X. 2018. Learning Discriminative Features with Multiple Granularities for Person Re-Identification. In *Proceedings of the 26th ACM international conference on Multimedia*, MM '18, 274–282. ACM.
- Wang, Z.; Fang, Z.; Wang, J.; and Yang, Y. 2020. Vitaa: Visual-textual attributes alignment in person search by natural language. In *European conference on computer vision*, 402–420. Springer.
- Wei, L.; Zhang, S.; Gao, W.; and Tian, Q. 2018. Person transfer gan to bridge domain gap for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 79–88.
- Wei, S.; Gao, Z.; Ma, C.; Zhao, Y.; Guan, W.; and Chen, S. 2025. Multiple Information Prompt Learning for Cloth-Changing Person Re-Identification. *IEEE Transactions on Image Processing*.
- Wu, X.; Ma, W.; Guo, D.; Zhou, T.; Zhao, S.; and Cai, Z. 2024. Text-based occluded person re-identification via multi-granularity contrastive consistency learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 6162–6170.
- Xia, J.; Tan, L.; Dai, P.; Zhao, M.; Wu, Y.; and Cao, L. 2023. Attention Disturbance and Dual-Path Constraint Network for Occluded Person Re-identification. *arXiv e-prints*, arXiv–2303.
- Yan, S.; Dong, N.; Zhang, L.; and Tang, J. 2023. Clip-driven fine-grained text-image person re-identification. *IEEE Transactions on Image Processing*, 32: 6032–6046.
- Yan, S.; Liu, J.; Dong, N.; Zhang, L.; and Tang, J. 2024. Prototypical prompting for text-to-image person re-identification. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 2331–2340.
- Yang, F.; Li, W.; Yang, M.; Liang, B.; and Zhang, J. 2024. Multi-modal disordered representation learning network for description-based person search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 16316–16324.
- Yuan, H.; Li, X.; Dai, J.; You, X.; Sun, Y.; and Ren, Z. 2025. Deep Streaming View Clustering. In *Forty-second International Conference on Machine Learning*.
- Zhao, Z.; Liu, B.; Lu, Y.; Chu, Q.; and Yu, N. 2024. Unifying multi-modal uncertainty modeling and semantic alignment for text-to-image person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 7534–7542.
- Zheng, L.; Shen, L.; Tian, L.; Wang, S.; Wang, J.; and Tian, Q. 2015. Scalable Person Re-Identification: A Benchmark. In *Proceedings of the IEEE International Conference on Computer Vision*, 1116–1124.
- Zheng, Y.; Zhao, Z.; Yu, X.; and Yu, D. 2022. Template-aware transformer for person reidentification. *Computational Intelligence and Neuroscience*, 2022(1): 8917964.
- Zhu, A.; Wang, Z.; Li, Y.; Wan, X.; Jin, J.; Wang, T.; Hu, F.; and Hua, G. 2021. Dssl: Deep surroundings-person separation learning for text-based person retrieval. In *Proceedings of the 29th ACM international conference on multimedia*, 209–217.
- Zhu, H.; Ke, W.; Li, D.; Liu, J.; Tian, L.; and Shan, Y. 2022. Dual cross-attention learning for fine-grained visual categorization and object re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4692–4702.