



Deposited via The University of York.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/239025/>

Version: Accepted Version

Article:

Alehdagh, Mahdi, Josi, Arthur, Cruz, Rafael et al. (2025) Adaptive Generation of Privileged Intermediate Information for Visible-Infrared Person Re-Identification. IEEE Transactions on Information Forensics and Security. ISSN: 1556-6013

<https://doi.org/10.1109/TIFS.2025.3541969>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Adaptive Generation of Privileged Intermediate Information for Visible-Infrared Person Re-Identification

Mahdi Alehdaghi, Arthur Josi, *Student Member, IEEE*, Rafael M. O. Cruz, Pourya Shamsolameli, *Senior Member, IEEE* and Eric Granger, *Member, IEEE*

Abstract—Visible-infrared person re-identification (V-I ReID) seeks to retrieve images of the same individual captured over a distributed network of RGB and IR sensors. Several V-I ReID approaches directly integrate the V and I modalities to represent images within a shared space. However, given the significant gap in the data distributions between V and I modalities, cross-modal V-I ReID remains challenging. A solution is to involve a privileged intermediate space to bridge between modalities, but in practice, such data is not available and requires selecting or creating effective mechanisms for informative intermediate domains. This paper introduces the Adaptive Generation of Privileged Intermediate Information (AGPI²) training approach to adapt and generate a virtual domain that bridges discriminative information between the V and I modalities. AGPI² enhances the training of a deep V-I ReID backbone by generating and then leveraging bridging privileged information without modifying the model in the inference phase. This information captures shared discriminative attributes that are not easily ascertainable for the model within individual V or I modalities. Towards this goal, a non-linear generative module is trained with adversarial objectives, transforming V attributes into intermediate spaces that also contain I features. This domain exhibits less domain shift relative to the I domain compared to the V domain. Meanwhile, the embedding module within AGPI² aims to extract discriminative modality-invariant features for both modalities by leveraging modality-free descriptors from generated images, making them a bridge between the main modalities. Experiments conducted on challenging V-I ReID datasets indicate that AGPI² consistently increases matching accuracy without additional computational resources during inference.

Index Terms—Visible-Infrared Person Re-Identification, Learning Under Privileged Information, Adaptive Image Generation.

I. INTRODUCTION

Person re-identification (ReID) involves matching images or videos of individuals of interest captured across multiple non-overlapping cameras. Current methods for person ReID, such as those based on deep Siamese networks [1]–[4], mainly use visible images to train a feature embedding backbone for similarity matching metric learning losses.

Although deep learning (DL) models have made significant progress in recent years, person ReID remains a challenging task due to the non-rigid structure of the human body and

variations in capture conditions in real-world scenarios. These variations can include illumination variations, blurring of motion, resolution, pose, viewpoint, occlusion, and background clutter [1], [5], [6]. Moreover, visible ReID systems cannot perform well in low-light environments (e.g., at night), thereby limiting their practical application.

Visible-infrared ReID (V-I ReID) is a variant of person ReID that involves matching individuals across RGB and IR cameras. V-I ReID is challenging since it requires matching individuals across different modalities with significant differences in appearance [7]–[10]. For robust cross-modal matching, a V-I ReID system must train a feature extractor that encodes pedestrians' images into feature information, which is interchangeable between V and I modalities. Then, during inference, images of people are matched across RGB and IR cameras based on these features. State-of-the-art methods for V-I ReID mainly focus on image-level [11], [12] or feature-level alignment [13]–[17]. Through the use of generative modules, the former aims to bridge the modality discrepancy by encoding images into feature space and then decoding and translating from different modalities to others. The latter obtains a modality-invariant embedding space by learning a shared model to represent images in feature space in which the discrepancy is minimized. Since V and I modalities have distinct data distributions, which results in a large discrepancy between V and I images in appearance and statistics [18], the features extracted from such models tend to be biased on modality-specific information and not be proper for cross-modal matching.

Addressing this challenge involves considering an intermediate space that provides meaningful information about the primary modalities, guiding the model to focus on shared knowledge. However, the availability of such intermediate spaces is not guaranteed, and their creation or identification is not a straightforward task. Additionally, determining how to effectively utilize bridging information during both the training and testing stages poses a key aspect of this solution.

The Learning Under Privileged Information (LUPI) paradigm [19] emerges as a solution aiming to leverage additional data available only during training. This approach has been adopted by various V-I ReID methodologies to bridge the gap between V and I images. For instance, works such as [20]–[25] concentrate on exploring subspaces between the IR spectrum and RGB channels, creating intermediate domains to address the modality gap. However, existing intermediate

M. Alehdaghi, A. Josi, R. Cruz, and E. Granger are with the LIVIA, ILLS, Dept. of Systems Engineering, ETS Montreal, Canada (Emails: {mahdi.alehdaghi.1, arthur.josi.1}@ens.etsmtl.ca, {Rafael.Menelau-Cruz, Eric.Granger}@etsmtl.ca).

P. Shamsolmoali is with the University of York, U.K. (Email: pshams55@gmail.com).

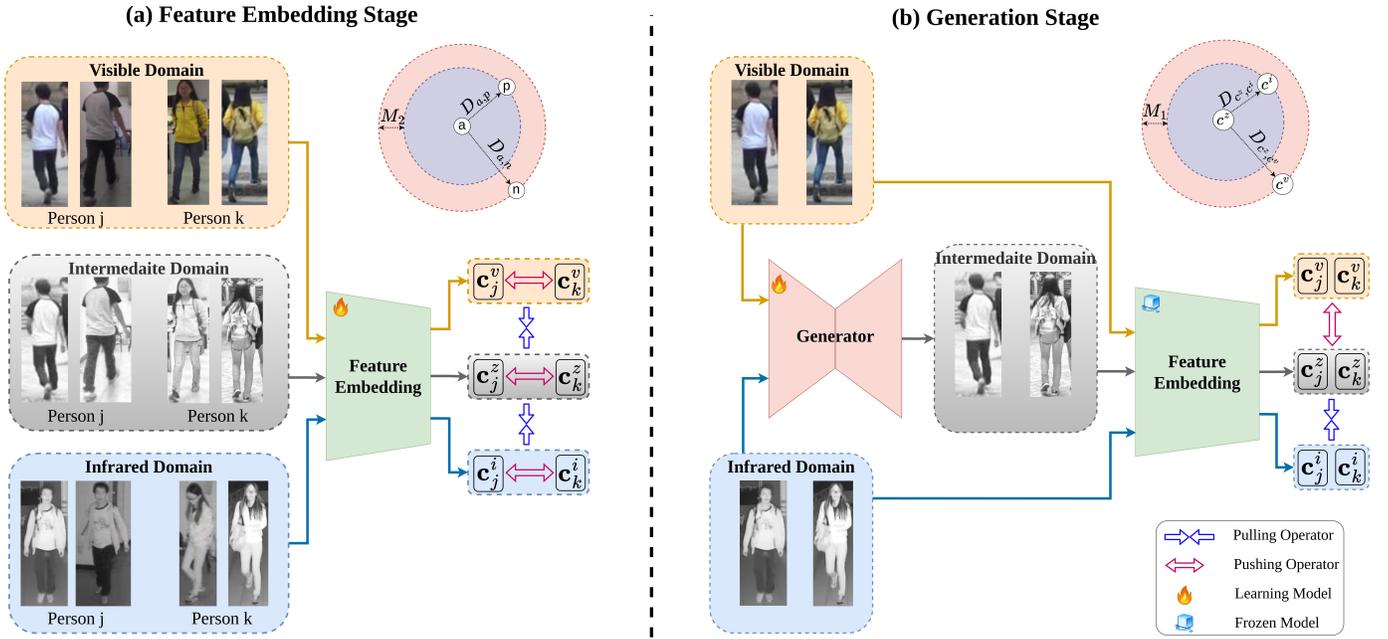


Fig. 1: A illustration of the training strategy with our AGPI² approach. The generated images are learned to be similar to the infrared modality by using adversarial objectives. (a) The feature embedding stage pushes the extracted features to approach the intermediate domain, while (b) the generation stage transforms V images to an intermediate domain that approaches I images. The objective of feature embedding is to minimize the distance from the anchor to the positive ($D_{a,p}$) and maximize its distance from the negative samples ($D_{a,n}$). Let c^v , c^z , and c^i be the center of the feature distributions for each class of V, intermediate, and I images. The generator tries to create a space in which the ID-aware features of each person are close to the infrared and visible at the same time. (To reach this goal and make the transferred images not be biased toward visible, we push them more toward infrared: $D_{c^z, c^i} < D_{c^z, c^v} + M_1$).

generative approaches fail to effectively adapt or generalize the intermediate information across V and I domains. The absence of specific objectives for generalizing intermediate images results in a limited ability to adapt to significant V-I domain shifts. To better capitalize on modality-invariant attributes, V-I ReID models necessitate appropriate intermediate images capable of bridging the substantial domain gap between I and V appearances. While utilizing intermediate images during testing enhances matching accuracy [11], [12], it introduces complexity and may not be ideal for real-time applications.

This paper introduces an approach for the Adaptive Generation of Privileged Intermediate Information (AGPI² approach) that trains a backbone model for V-I person ReID by adapting a new intermediate space that connects V and I modalities. Our proposed AGPI² training approach comprises three key modules: a feature embedding backbone, a generator, and an ID-modality discriminator. Since each modality can be seen as a domain with its data distribution, we propose treating each modality as a separate domain. Then, our AGPI² aim is to generate an intermediate domain with similar attributes to the source V and I domains. This intermediate space enables the model to learn and extract common attributes between the two modalities. At inference time, the model can match the same person across modalities by representing them in a non-modality-specific feature space without relying on intermediate information. To achieve this, our framework adapts intermediate images to have the same content (body

pose, background, and texture) as V images while sharing the style information (heat or temperature information) with I images. The AGPI² component provides an intermediate step for the feature embedding backbone, helping to reduce modality-specific information in the extracted features.

Given the absence of supervised information to generate the intermediate domain containing ID-aware information and bridging modality attributes, we propose an end-to-end approach with adversarial objectives w.r.t. the ultimate goal of V-I ReID, which is extracting unique and modality-free descriptors for each person. To this end, the generative module is encouraged to transform the V images to an intermediate domain to minimize the domain gap between such virtual and I images. Consequently, the generative module must focus on ID-aware attributes to avoid the intermediate images missing the discriminative information about the person in the V images. To reach this goal, the ID-modality discriminator is proposed, which seeks to reinforce the generator by discriminating modalities from ID-related features. Such property in the discriminator forces the generator to focus on ID-aware information. The feature embedding module simultaneously works to minimize the gap between the generated and V images, as shown in the left side of Fig. 1, while the generator's objective is to maximize this distance (right side of 1). This collective approach ensures the effective generation of intermediate domain images that capture both ID-aware information and modality attributes.

In other words, AGPI² tackles the significant gap between modalities by generating and utilizing bridging information that links the two domains. Since such information is not available all the time or hard to benefit from, our AGPI² training approach is on two stages: (a) Finding a privileged linking domain between two modalities and (b) making the model robust against modality change by leveraging the intermediate domain as a clue for containing shared attributes among two modalities. In comparison with [26], [27] that are creating an intermediate domain in the latent feature space, our intermediate domain is in the input space. Therefore, our generative module can be used in other applications to generate intermediate images between visible and infrared and even in other V-I person ReID. In Fig. 1, the generative module is depicted as trained to generate images with features closer to I and farther from V, enforcing a margin M_1 between the distances from the center of the feature vector representing images of each person in the intermediate space to V (D_{c^z, c^v}) and I (D_{c^z, c^i}). Simultaneously, the feature embedding stage ensures close similarity between the feature vectors of each person in the visible, intermediate, and infrared modalities.

Our main contributions are summarized as follows: (1) An adversarial training strategy called AGPI² is proposed to reduce the V and I domain gaps by adopting a generalized intermediate domain as privileged information. (2) By analyzing the mutual information between intermediate space and V-I person ReID objectives, adversarial losses are proposed between the ID-modality discrimination, feature embedding, and generation modules to encourage the transformed images to incorporate discriminative ID information while reducing the influence of domain-specific attributes. Such losses make the modules play a min-max game between each other, resulting in a balanced and informative intermediate space. (3) An extensive set of experiments on the challenging SYSU-MM01 [28] and RegDB [29] datasets indicate that our proposed AGPI² not only outperforms state-of-art methods for cross-modal V-I person ReID but can also be seamlessly integrated into other methods to enhance their accuracy without imposing additional computational overhead during inference. In other words, our framework is agnostic and can be applied to other image retrieval applications to improve their performance.

II. RELATED WORK

(a) Visible-Infrared Person ReID:

This paper focuses on V-I ReID, where deep backbone models are trained to match individuals between V and I cameras using a labeled set of V and I images captured using RGB and IR cameras, respectively. Then, during inference, different individuals are matched across camera modalities, that is, matching query I images to gallery V images, or vice versa [16], [17], [30]–[35]. Existing DL models used for V-I ReID can be divided into Generative and Representative models. The former seeks to map raw data between two modalities to reduce modality discrepancies at the image level, while the latter focuses on finding a suitable discriminant representation by reducing the cross-modality gap at the feature level.

Some generative models, like ThermalGAN [12] or AlignGAN [11] employed autoencoders to translate V and I images

and found a common feature space through the generation process. D²RL [36] generated fake-real paired images from synthetic V and I images and encoded them into a common feature space for feature representation and matching. Other models [37]–[39] relied on GANs to create synthetic images conditioned by a modality-invariant representation. However, these methods are complex at inference time, requiring image generation at test time. Moreover, they do not necessarily produce high-quality images that are capable of bridging the gap between modalities.

Representation approaches focus on training a backbone model to learn a discriminant representation that captures the essential features of each modality while exploiting their shared information [13]–[15], [31], [35], [40] or exploiting self-similarity [41]. In [13], [17], [23], [42]–[44], authors only used global features combined through multimodal fusion. The absence of local information limits their accuracy in more challenging cases. To improve robustness, recent methods combined global information with local part-based features [10], [45]–[50] or fused features of different resolutions similar to LRAR [51], which used two resolution-adaptive mechanisms. For example, [45], [46] divided the spatial feature map into fixed horizontal sections and applied a weighted-part aggregation. Using the weighted mechanism, these methods could dynamically pay attention to the relative importance of different parts to improve their discrimination power. [48], [49], [52] dynamically detected regions in the spatial feature map to address the misalignment of fixed horizontally divided body parts.

In contrast, the proposed AGPI² uses a generative model only during the training phase to generate the privileged information to help the training process by bridging the modality gap. This makes our proposed method computationally efficient, relying only on the feature embedding backbone during inference.

(b) Intermediate Modality in V-I ReID: Given the significant gap between V and I distributions, cross-modal matching is a challenging problem. One possible approach is to use an intermediate modality to help the model find common semantics between these two modalities. For example, the model [40] relied on gray images instead of colored images. However, using grayscale alone limited the amount of valuable information from V images by not containing the V spectrum information. Therefore, Fan et al. [53] combined the information from the three channels (V, I, and grayscale) and [23] used grayscale images as an intermediate modality during the training process, while in [42], [54], [55] augmented images were created by randomly selecting one value from the RGB color as the intermediate modality. Alehdaghi et al. [20] introduced a novel approach involving the use of random linear combinations of RGB colors. This technique served as an effective bridge, facilitating the extraction of color-independent features and ensuring that the derived attributes are not only devoid of color bias but also contain characteristics commonly found in infrared images. [55] performed the channel-level interaction by randomly determining whether to keep the original RGB image or perform a random channel selection. However, these methods made limited use of intermediate

domains by searching for them in a small search space.

To address this limitation, methods in [7], [18], [21], [24] used generative models to create the wider search space for intermediate modality. For example, [21] used a single convolution layer to transform V images into a new modality without using the I or identity information. [18] proposed a pixel-to-pixel feature fusion operation on V and I images to build the synthetic images, and [24] introduced a shallow auto-encoder to generate intermediate images from both modalities. The images generated by these methods were biased toward the source modality. Lu et al. [7] designed a channel interactive generator to generate confused modality images to address this problem. While such methods enhance performance, they ignore the intricate V-I image relationship and fail to take into account domain gaps between id-discriminative information during the generation of intermediate domains. In contrast, our approach focuses on minimizing domain shifts between visible and infrared to identify the most informative intermediate space. It also employs an ID-modality discriminator to focus on the person's identity, allowing the generative model to robustly create persons' images from one modality into an intermediate while preserving the discriminative attributes of that person. Thus, the quality of the intermediate image and the id-aware information is significantly improved.

(c) Learning Under Privileged Information: Supervised learning is based on labeled data to optimize an appropriate decision function, minimizing the generalization error. The learner's function typically converges rapidly to the optimal solution for simple learning tasks with a large amount of training data. However, the convergence rate becomes slow for more challenging tasks and larger optimization spaces [56]. The LUPI approach was initially proposed to accelerate the convergence of SVMs by incorporating prior knowledge into the learning process. In [57] used a hallucination network to introduce prior knowledge into an additional stream in CNN to extract depth-related attributes. Beyond seeking faster convergence, LUPI has been successfully applied in various applications to improve performance [58]–[62]. In [63], a GAN-based model is proposed to generate prior knowledge at the testing time, and Crasto et al. [64] distilled this knowledge from privileged information to the V branch. Although the use of PI improves model performance, such proper and supervised information is unavailable or difficult to leverage. In this study, our focus lies in implementing the LUPI paradigm for V-I ReID, aiming to identify suitable and informative Privileged Information (PI) during training. To achieve this objective, we propose a generative model designed to dynamically transform modalities into an intermediate space, serving as privileged information exclusively during the training phase through the use of adversarial losses.

III. PROPOSED METHOD

Our AGPI² approach generates a virtual PI domain that bridges V and I modalities during training and then leverages it to learn modality-invariant features. Fig. 2(a) shows the overall AGPI² training architecture, comprising three modules: (1) The generator module that transforms V images into

intermediate ones by incorporating infrared features. These intermediate images help the model simulate an infrared version of the visible images, allowing it to focus on non-color-specific attributes from the visible modality. Since such images have exact pose information of individuals and only lose the colors, the model accesses one intermediate step before seeing infrared images with larger shifts (different pose, background, and modality). (2) ID-modality discriminator that identifies modality from ID-aware features to avoid the intermediate images being over-focused on the background, which is likely to be dominant for modality transformation, and (3) A feature embedding module that extracts shared features from V and I images, guided by the proposed bridging losses to effectively minimize modality discrepancies.

During training, the ID-modality discriminator is taught to classify individuals in the intermediate and V images, considering them as belonging to the same class but distinct from the I images. This is achieved through the use of identity discriminative features extracted by the backbone. The primary goal is to guide the generative module in prioritizing the capture of personal attributes aligned with the characteristics associated with modality-specific information. Concurrently, in the generation learning stage, the generator aims to produce intermediate images that the discriminator would adversarially classify as belonging to the I modality. This iterative process leads to the generated images becoming more similar to the I images, thereby enhancing the model's ability to extract common features from both V and I modalities.

Preliminaries: Let us assume a multimodal dataset for V-I ReID problems that is represented by $\mathcal{S} = \{\mathcal{V}, \mathcal{I}\}$, where $\mathcal{V} = \{v_j\}_1^n$ are V and $\mathcal{I} = \{i_j\}_{n+1}^{n+m}$ are I images of N different persons, and $\mathcal{Y} = \{y_j\}_1^{n+m}$ contains their identity for the training process. Metric losses \mathcal{L} are typically defined based on distances D between feature vectors $\mathbf{f}_j^v = \varphi(v_j; \Theta)$ of V images and feature vectors $\mathbf{f}_j^i = \varphi(i_j; \Theta)$ of I images for every two distinct individuals, as:

$$\mathcal{L} = \sum_{j=1}^{n+m} \mathcal{C}(D(\mathbf{f}_j^v, \mathbf{f}_j^i), D(\mathbf{f}_j^v, \mathbf{f}_k^v), y_j, y_k), \quad (1)$$

where \mathcal{C} is metric loss objectives which can be triplet-loss [65] or contrastive loss [66], y_j and y_k represent the identity of person in images j and k , and $\mathbf{f}_j^v, \mathbf{f}_k^i \in \mathbb{R}^d$, where d the features dimension.

A. LUPI Framework

Given that the generated images are unavailable at test time, we formulate our proposed learning strategy according to the LUPI paradigm. With LUPI [67], the DL model leverages privileged information (PI) exclusively during the training phase, while the main information remains accessible in both the training and testing phases. For example, in cross-modal matching, the model typically relies on the V and I modalities during the training phase to leverage shared discriminative features between inputs. However, during the inference phase, the model only inputs the query image from one modality (V or I). Also, during the training phase, the model can use privileged bridging information between the two modalities to

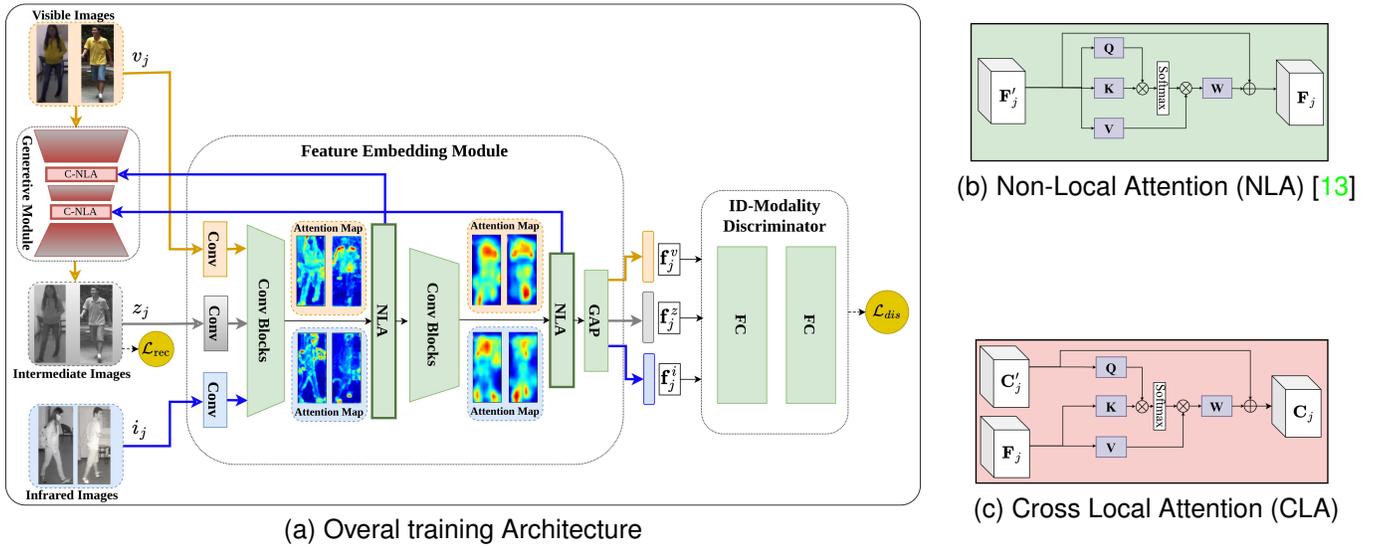


Fig. 2: Block diagram of our proposed AGPI² training strategy for V-I person ReID. It takes advantage of the adaptive intermediate domain to reduce the distribution gap between the V and I domains. The feature embedding module seeks to minimize the distance between the two original (I and V) modalities by using the auxiliary generated intermediate domain. The ID-modality discriminator learns to detect the modality from ID-aware features, and the generation process aims to transform V images to infrared through adversarial training with the discriminator.

address the significant domain discrepancies between V and I images [56]. By incorporating the intermediate modality (Z), AGPI² reduces the modality gap between V and I, enabling the model to generalize more effectively and improve its ability to transfer knowledge across both modalities.

B. Mutual Information Analysis

We formulate the bridging characteristics objectives of the privileged intermediate information from the mutual information perspective, and demonstrate that jointly learning id-aware and modality-free objectives achieves a conditional mutual information maximization between intermediate images and identity space while discarding the modality style information:

$$\max_Z \text{MI}(Z; Y|M), \quad (2)$$

where Y is the identity of person input images, M is the modality label, and $\text{MI}(\cdot; \cdot)$ is the mutual information between two random variable. For solving Eq. (2), we can expand it as follows:

$$\text{MI}(Z; Y|M) = \overbrace{\text{MI}(Z; Y)}^{\text{ID-Aware}} - \overbrace{\text{MI}(Z; Y; M)}^{\text{ID-Modality-Aware}}. \quad (3)$$

To maximize this, the first term (ID-aware) should be maximized and the second term (ID-Modality-Aware) minimized. To minimize the ID-Modality-Aware part for intermediate images, an adversarial training approach is proposed between the ID-Modality Discriminator and the generative module, in which the former maximizes that part by looking for id-related information specific to the modality. At the same time, the latter tries to create Z images that fool the discriminator, thereby minimizing this term.

C. ID-Modality Discriminator

For the generator to create intermediate images that contain id-related information, we propose a module to discriminate modality along with the identities of each individual in images. To achieve this, we employ an MLP network as our ID-modality Discriminator to classify modality types (Infrared or Visible) from features extracted from the shared embedding backbone. The generator's objective is to deceive the discriminator by creating a modality that appears false for a specific identity.

To implement this idea, we doubled the label space of each identity to train the discriminator to differentiate between the same identity in the V and I modalities:

$$y'_j = \begin{cases} 2y_j & \text{V and Intermediate modalities} \\ 2y_j + 1 & \text{I modality} \end{cases} \quad (4)$$

where y_j is the original label.

For training the ID-modality discriminator, the objective is defined as:

$$\mathcal{L}_{\text{dis}} = \sum_{l=1}^{2N} -y'_{j,l} \log(p'_{j,l}), \quad (5)$$

where $p'_{j,c}$ is the predicted probability observation j of class c by the discriminator \mathcal{D} .

Fig. 3 shows the main difference between the general modality discriminator and our ID-modality discriminator. Each training sample is a shape labeled with the target label. Other SOTA intermediate generative approaches use ordinary binary discriminators that are trained to identify only modality (V and intermediate as the same), disregarding the person's identity. In contrast, by proposing new label spaces that are related to the identities of the discriminator, it would be able to detect V/I modalities from the ID-aware attributes of samples during training.

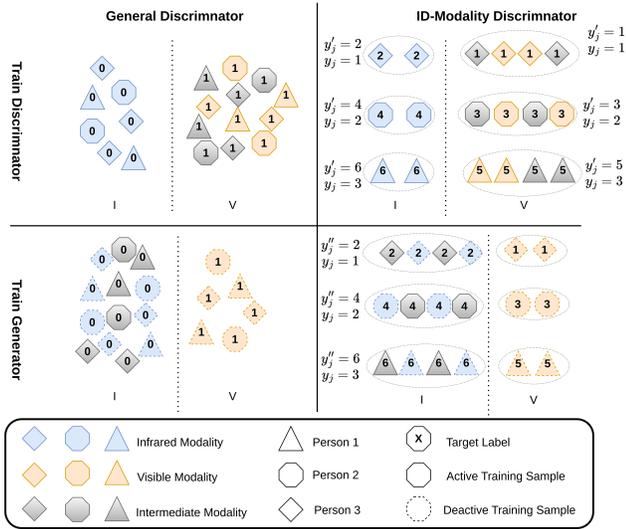


Fig. 3: ID-Modality discriminator vs. general discriminator. Our label space is doubled to account for identity (differentiate between individuals in each modality).

D. Generative Module

To generate the intermediate images, an adaptive and unsupervised intermediate generative model \mathcal{G} is proposed to transfer visible images to intermediate modality in an end-to-end fashion regarding Eq. (2). The reason behind using visible images is that the visible images contain more modality-specific id-discriminative information [68] and need to be filtered to remove this biased information. To have information from both modalities in the generation process, the V images, at first, are encoded as features vectors and fused with I features F_j^i by the cross non-local attention [69], and then the decoder reconstructs the image, $z_j = \mathcal{G}(v_i, F_j^i)$. To allow this module to focus on ID-aware information in intermediate images and maximize the $\text{MI}(Z; Y)$, the cross-entropy between identity space and intermediate images must be minimized (we proved this in Appendix A-A). So the \mathcal{L}_{id} which is a cross-entropy loss function between the features extracted from intermediate images (\mathbf{f}_j^z) and identity label y_j is proposed to be minimized. Also, to filter them from any modality-specific information, these images should maximize the \mathcal{L}_{dis} when they are fed to the ID-modality discriminator. To make this, we need a new label to fool the discriminator, and a label that specifies the same identity but in a different modality is chosen:

$$y'_j = 2y_j + 1. \quad (6)$$

The reason is that since this discriminator is not binary, every label except the ground truth with which the discriminator is trained can be selected as an opposite label. So, the proper label should be selected to avoid interrupting the id-discriminative information in the created images.

Also, since the Z images are created by transferring from V to intermediate space, they tend to be biased toward visible space. To tackle this issue, we pushed their feature more

Algorithm 1 Joint AGPI² Training Strategy of \mathcal{F} , \mathcal{G} and \mathcal{D}

Require: $\mathcal{S} = \{\mathcal{V}, \mathcal{I}\}$ as training data
while all batches are not selected **do**
 $v_j, i_j \leftarrow \text{batchSampler}(b, p)$
 $i_j^* \leftarrow \mathcal{G}(i_j, F_j^i)$ \triangleright reconstruct I
 $z_j^* \leftarrow \mathcal{G}(v_j, F_j^i)$ \triangleright generate intermediate with detached gradient style feature F_j^i
 unfreeze \mathcal{F} and \mathcal{D}
 update $\Theta_{\mathcal{F}}$ by optimizing Eq. (12) \triangleright detach gradient from z_j to not modify \mathcal{G}
 update $\Theta_{\mathcal{D}}$ by optimizing Eq. (5) while freezing \mathcal{F}
 update $\Theta_{\mathcal{G}}$ by optimizing Eq. (9) while freezing \mathcal{D}
end while

toward infrared than visible features. This adversarial objective can be defined as:

$$\mathcal{L}_{adv} = \sum_{l=1}^{2N} -y''_{j,l} \log(p''_{j,l}) + \max(0, M_1 + D_{\mathbf{c}_j^z, \mathbf{c}_j^i} - D_{\mathbf{c}_j^z, \mathbf{c}_j^v}), \quad (7)$$

where p''_j is the predicted label j from the discriminator, and $\mathbf{c}_j^z, \mathbf{c}_j^i, \mathbf{c}_j^v$ are the centers of feature embedding of generated I, and V images, respectively, for the person with identity y_j . The function D measures the distance between two vectors.

To generate more realistic images, we also transfer I images into the intermediate domain, which provides a supervised reconstruction loss:

$$\mathcal{L}_{rec} = |\mathcal{G}(i_j, F_j^i) - i_j|, \quad (8)$$

and the total loss for the generative modules is:

$$\mathcal{L}_{gan} = \mathcal{L}_{rec} + \mathcal{L}_{idz} + \lambda_{adv} \mathcal{L}_{adv}. \quad (9)$$

E. Feature Embedding Module

The feature embedding module, \mathcal{F} , has three backbones with shared layers to extract features from V, intermediate, and I images. Modified versions of the triplet and color-free losses are proposed to encourage modality-invariant descriptor features and train the V-I ReID model while considering intermediate images.

(1) Color-Free Loss: Previous studies [23], [56] have shown the advantages of feature invariance when extracting features from both original V and generated modality images Z. For feature representations that are robust to modality variations, we used the same color-independent loss over virtual and V images as defined in [56] for training the V and intermediate backbone:

$$\mathcal{L}_{cf} = \|\mathbf{f}_j^v - \mathbf{f}_j^z\|, \quad (10)$$

where \mathbf{f}_j^z is the feature vector of the intermediate images z_j^* .

(2) Intermediate Dual Triplet Loss: Since the role of intermediate space is bridging between modalities, we use extracted features from intermediate images as the middle between them and apply two separate triplet losses in dual mode:

$$\mathcal{L}_{dual} = \mathcal{L}_{tri}(\mathbf{f}_{a \in \mathcal{V}}, \mathbf{f}_{p \in \mathcal{I}}, \mathbf{f}_{n \in \mathcal{Z}}) + \mathcal{L}_{tri}(\mathbf{f}_{a \in \mathcal{I}}, \mathbf{f}_{p \in \mathcal{Z}}, \mathbf{f}_{n \in \mathcal{V}}), \quad (11)$$

where \mathcal{L}_{tri} [13], [45] is common soft-margin triplet loss [13], [45] with margin M_2 , a , p , and n indicate the anchor, positive,

Family			SYSU-MM01						RegDB					
			All Search			Indoor Search			Visible → Infrared			Infrared → Visible		
Method	Venue		R1	R10	mAP	R1	R10	mAP	R1	R10	mAP	R1	R10	mAP
Generative	AGAN [11]	ICCV19	42.4	85.0	40.7	45.9	92.7	45.3	57.9	-	53.6	56.3	-	53.4
	D ² RL [36]	CVPR19	28.9	70.6	29.2	-	-	-	43.4	66.1	44.1	-	-	-
	JSIA [38]	AAAI20	45.01	85.7	29.5	43.8	86.2	52.9	48.5	-	49.3	48.1	-	48.9
	Hi-CMD [37]	CVPR20	34.94	77.58	35.94	-	-	-	86.39	-	66.04	-	-	-
	TS-GAN [70]	PRL21	58.3	87.8	55.1	62.1	90.8	71.3	68.2	-	69.4	-	-	-
Representation	AGW [13]	TPAMI20	47.50	-	47.65	54.17	-	62.97	70.05	-	50.19	70.49	87.21	65.90
	DDAG [45]	ECCV20	54.74	90.39	53.02	61.02	94.06	67.98	69.34	86.19	63.46	68.06	85.15	61.80
	SSFT [46]	CVPR20	63.4	91.2	62.0	70.50	94.90	72.60	71.0	-	71.7	-	-	-
	SIM [68]	IJCAI20	56.93	-	60.88	-	-	-	74.47	-	75.29	75.24	-	78.30
	MPANet [48]	CVPR22	70.58	96.21	68.24	76.74	98.21	80.95	83.70	-	80.9	82.8	-	80.7
	DTRM [55]	TIFS22	63.03	93.82	58.63	66.35	95.58	71.76	79.09	92.25	70.09	78.02	91.75	69.56
	FTMI [71]	MVA23	60.5	90.5	57.3	-	-	-	79.00	91.10	73.60	78.8	91.3	73.7
	RAPV-T [72]	ESA23	63.97	95.30	62.33	69.00	97.39	75.41	86.81	95.81	81.02	86.60	96.14	80.52
Intermediate	HAT [23]	TIFS20	55.29	92.14	53.89	-	-	-	71.83	87.16	67.56	70.02	86.45	66.30
	CAJ [42]	ICCV21	69.88	-	66.89	76.26	97.88	80.37	85.03	95.49	79.14	84.75	95.33	77.82
	RPIG [20]	ECCVw22	71.08	96.42	67.56	82.35	98.30	82.73	87.95	98.3	82.73	86.80	96.02	81.26
	MMN [24]	ICM21	70.60	96.20	66.90	76.20	99.30	79.60	91.60	97.70	84.10	87.50	96.00	80.50
	SMCL [22]	ICCV21	67.39	92.84	61.78	68.84	96.55	75.56	83.93	-	79.83	83.05	-	78.57
	G2DA [73]	PR23	63.94	93.34	60.73	71.06	97.31	76.01	-	-	-	-	-	-
	MCBD [7]	TIFS23	71.63	94.98	67.28	79.44	98.32	79.85	90.10	97.09	82.73	87.66	96.68	81.52
	AGPI ² (ours)	-		72.23	97.04	70.58	83.45	98.62	84.25	89.03	98.19	83.89	87.91	97.15

TABLE I: Accuracy of the proposed AGPI² and state-of-the-art methods on the SYSU-MM01 (single-shot setting) and RegDB datasets. All numbers are percent.

Family		LLCM			
		V → I		I → V	
Method	Venue	R1	mAP	R1	mAP
DDAG [45]	ECCV20	40.3	48.4	48.0	52.3
CAJ [42]	ICCV21	56.5	59.8	48.8	56.6
RPIG [20]	ECCVw22	57.8	61.1	50.5	58.2
AGPI ² (ours)	-	61.78	65.08	56.47	62.79

TABLE II: Accuracy of the proposed AGPI² and state-of-the-art methods on the Large LLCM Dataset. All numbers are percent.

and negative samples, respectively. The superscripts i , v , and z select V, I, and intermediate features, respectively.

The total loss proposed to train the feature embedding backbone is:

$$\mathcal{L}_f = \mathcal{L}_{id} + \mathcal{L}_{dual} + \lambda_{cf}\mathcal{L}_{cf}. \quad (12)$$

where \mathcal{L}_{id} is the cross-entropy loss to allow identity-discriminative features [13].

Joint Learning: Each batch of data contains b person with p positive images from the V and I domains. After extracting the embedding features, the model generates intermediate images as described in Algorithm 1. Our joint training strategy optimizes the generator \mathcal{G} by freezing the \mathcal{F} and \mathcal{D} parameters and also detaches the gradients from the generated images when the parameters \mathcal{F} are updated.

IV. RESULTS AND DISCUSSION

A. Experimental Methodology

Datasets: Cross-modal research has made extensive use of the challenging SYSU-MM01 [28], RegDB [29] and recently

published LLCM [74] datasets. SYSU-MM01 is a large-scale dataset that contains more than 22K and 11K V, and I images that are not co-located belong to 491 individuals, respectively. These images were captured from four RGB and two near-infrared cameras, with 395 identities used for training and 96 for testing. The dataset has two evaluation modes, single-shot and multi-shot, based on the number of images in the gallery. The former considers only one random image per identity in the gallery, while the latter considers ten images per identity. The RegDB dataset comprises 4,120 co-located V-I images from 412 identities captured from a single camera. Ten trial configurations randomly divide the dataset into two identical sets (206 identities) for training and testing. Tests are conducted in two different ways, comparing I to V (query) and vice versa. An LLCM dataset consists of a large, low-light, cross-modality dataset that is divided into training and testing sets at a 2:1 ratio.

Implementation Details: Pre-trained ResNet50 [75] with two non-local attention modules is used as our feature extractor, the same as the AGW model [13]. Also, concerning the analysis of different models and inputs in IV-F, we choose VQ-VAE [76] as the adaptor for V images into the intermediate domain. All inputs are resized to 288 by 144, then cropped randomly and filled with zero padding or mean pixels. SGD and ADAM optimizers were used for the optimization process of re-identification and generative modules, respectively. To perform the triplet loss, a minimum of two distinct individuals must be included in each mini-batch. So, we set $p = 4$ different images of $b = 8$ distinct persons from the dataset for each batch. $M_1 = 0.1$, $M_2 = 0.3$, $\lambda_{adv} = 0.1$, and $\lambda_{cf} = 10$ were set based on the analysis conducted in B-A.

Performance Measures: We use Cumulative Matching Char-

acteristics (CMC), and Mean Average Precision (mAP) as evaluation metrics. The CMC is rank-k accuracy, which determines the probability of a correct cross-modality person image appearing in the top-k retrieved results. In contrast, when multiple matching images are found in a gallery, mAP measures the image retrieval performance.

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

Table I provides a comparison of our proposed AGPI² with state-of-the-art V-I ReID approaches. Our experiments show AGPI² outperforms these methods in various situations. Although our approach requires additional image generation during training, this process is not required during inference, making it suitable for real-time applications. Additionally, the model learns to robustly represent features across modalities. In particular, our model improves the R1 accuracy and mAP score on the large-scale SYSU-MM01 dataset by 24.73% and 22.93%, over AGW’s [13] (our baseline). Compared to the second-best approach for the “indoor search” scenario, AGPI² outperforms by a margin of 1.15% R1 and 2.34% mAP without adding complexity to the feature backbone. For the RegDB dataset, AGPI² outperforms by a margin of R1 and mAP by 1.08% and 1.16%, for “visible to thermal” ReID scenarios. Similar improvements are noticeable in the “thermal to visible” scenario. In addition, our approach has the advantage that it can be easily used in different cross-modal ReID models to mine PI only in the training phase without incurring overhead during testing.

The AGPI² model achieves improvements on the large and complex LLCM dataset as shown in Table II. Since this dataset was introduced recently, few papers reported their results, so we ran other competitors’ approaches with published code on the LLCM dataset.

C. Incorporating to Other State-of-the-arts Models

The effectiveness of AGPI² is evident when integrated into state-of-the-art V-I ReID models [45], [48], [77]. We executed the authors’ provided code for each method, using their specified hyperparameters, both with and without AGPI². Table III highlights that AGPI² significantly improves performance when incorporated into these baseline models. For instance, integrating AGPI² into DDAG [45] and SEFL [77] boosts mAP accuracy by 4% and 1%, respectively. Notably, SEFL uses supervised body-shape information from modalities and also utilizes gray images as a bridging modality alongside visible and infrared images in training, while DDAG is trained solely on visible and infrared images. Therefore, our AGPI² framework showcases its bridging ability more effectively in DDAG compared to SEFL.

D. Domain Shift Between Modalities

As we discussed earlier, the privileged intermediate domain bridges the main domains by sharing discriminative attributes that are common to both. Crucially, this intermediate domain must remain unbiased towards any particular domain and lie at a balanced point. To validate that our created virtual domain meets this criterion, we measured the MMD distance between

Method	R1 (%)	mAP (%)
DDAG [45]	52.74	51.92
DDAG with AGPI ²	57.38	55.27
MPANet* [48]	70.58	68.24
MPANet* with AGPI ²	72.85	70.66
SEFL [77]	74.05	69.21
SEFL with AGPI ²	75.17	71.43

TABLE III: Accuracy of SOTA methods with AGPI² on the SYSU-MM01, testing under single-shot setting. Results were obtained by executing the author’s code on PyTorch v1.12 on Ubuntu 22.04 with 40 GB NVIDIA A100 GPU.

Modality	SYSU-MM01		RegDB	
	I	V	I	V
Infrared	0	0.86	0	1.05
Visual	0.86	0	1.05	0
Grayscale (HAT [23])	0.68	0.31	0.93	0.36
Random Comb. (RPIG [20])	0.64	0.34	0.92	0.39
MMN [24]	0.61	0.33	0.87	0.43
AGPI ² (ours)	0.58	0.43	0.79	0.52

TABLE IV: MMD between V, I, and different intermediate generation images on the SYSU-MM01 and RegDB datasets. For the I column, lower is better, and for the V, higher is better.

the intermediate domain and both V and I domains that reflect a balanced intermediary position, ensuring an unbiased connection.

The kernel-based Maximum Mean Discrepancy (MMD) [78] is used to measure the distance between deep features extracted by a ResNet50 model trained on ImageNet [79]. Table IV shows this distance for V, I, and intermediate images for the SYSU-MM01 dataset. As described in Section III, the intermediate domain must bridge the main domains to each other. Hence, it needs to have a smaller discrepancy to V and I, i.e., $MMD(V, Z) \leq MMD(I, V)$ and $MMD(I, Z) \leq MMD(I, V)$. For example, on the SYSU-MM01 dataset, the discrepancy between the intermediate images generated using our AGPI² method and I is 0.58, which is lower than the V-I discrepancy of 0.86. When compared to fixed gray images [23], which remain static and fail to adapt to the infrared spectrum, random gray images [20] exhibit superior bridging capabilities. This advantage stems from their lack of bias towards specific color spectrums. In contrast, AGPI² images strike a balance between I and V modalities, leveraging adaptive information across color spectrums during generation. This characteristic ensures that the intermediate space provides optimal information for the feature backbone to leverage common discriminative features between V and I images. Also, for visual comparison, we show examples of images created by our approach in Fig. 4. The third column shows intermediate images that share the same pose as the first column (V), seek to incorporate the style of the second column (I), and lose the color information.

E. Ablation Studies

A thorough ablation study is conducted to assess the impact of each module on AGPI² performance through various exper-

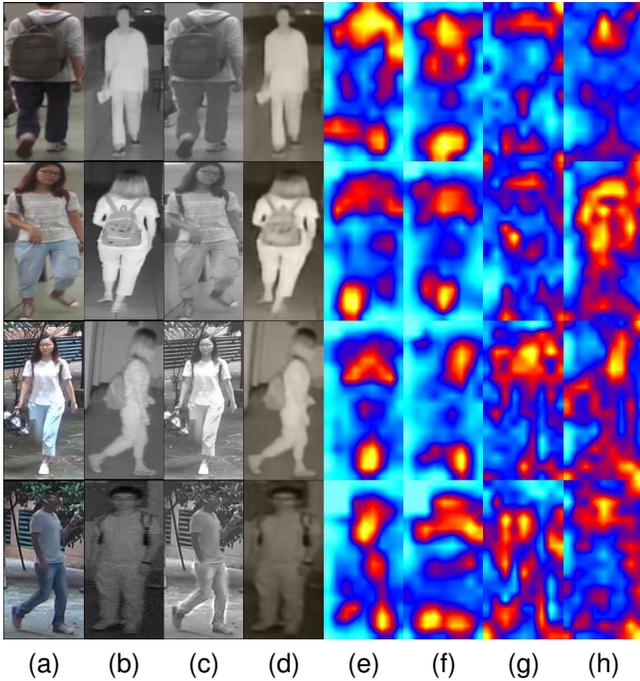


Fig. 4: Examples of images generated with our AGPI² for SYSU-MM01 dataset. Columns (a) and (b) are V and I images. The intermediate and reconstructed I images are shown in (c) and (d), respectively. The GradCAM [80] of our ID-Modality vs binary discriminator is shown in columns (e) to (h). Columns (e) and (f) for AGPI² discriminator for V and I images, respectively, and (g),(h) for Binary discriminator.

iments on the SYSU-MM01 dataset. Initially, we compare our intermediate domain approach with other methods that use fixed and generated intermediates. Subsequently, to validate the roles of loss objectives, the generative module, and the ID-modality discriminator, we experiment with different settings and combinations with other methods.

1) *Effect of Intermediate Image Generation:* Table V shows the performance of cross-modality ReID with and without using an intermediate module. With the adaptive generated images (I-V-Z training), we achieve 56.35%, while using the baseline model (I-V), we get 41.57% mAP. Additionally, the adaptive version shows comparable performance in both V and I testing settings, making it an effective alternative for a robust person ReID. Table VI provides different options for intermediate images and it is shown that, in comparison to other approaches using the intermediate domain, AGPI² has better performance.

2) *Effect of Losses:* In the following, we discuss the effect of each loss in AGPI².

Effect of Intermediate Dual Triplet Loss: Table VII shows that when used in conjunction with the dual triplet loss, this loss provides supervision to improve performance at test time. Although the ordinary triplet loss improves performance when it is integrated with \mathcal{L}_{cf} , it ignores the privileged information provided by the intermediate domain. To show that fact, we compared the effect of our \mathcal{L}_{dual} to \mathcal{L}_{tri} by using each of them

Training	Testing		R1 (%)	mAP(%)
	Query	Gallery		
I-V	V	V	97.40	91.82
	I	I	95.96	80.49
	I	V	58.69	41.57
I-V-Z	V	V	97.68	90.19
	I	I	95.34	81.62
	I	V	73.17	56.35

TABLE V: The impact on the accuracy of the proposed intermediate domain in SYSU-MM01 under the multi-shot setting.

Model	R1(%)		mAP(%)	
	S	M	S	M
Pre-trained (test: I→V)	2.26	3.31	3.61	1.71
Lower (train: V, test: I→V) *	7.41	9.86	8.83	6.83
Upper (train: I, test: I→I) §	90.34	94.90	87.45	80.38
Baseline (test: I→V)	49.41	58.69	41.55	38.17
Grayscale [23] (test: I→V)	61.45	65.16	59.46	46.94
RandG [20] (test: I→V)	64.50	69.01	61.05	50.84
MUN ⁺ [26] (test: I→V)	71.66	-	68.35	-
AGPI ² (test: I→V)	72.23	78.19	70.58	64.13

TABLE VI: Accuracy using different intermediate images on SYSU-MM01 using the Single and Multi-shot gallery. "*" indicates that our model is trained on V images while tested with I. In "§," we use I samples to train and test the model. "+" means that we only reported the MUN [26] model performance of using auxiliary bringing part.

alongside other losses in the training process. As can be seen, using such regularization increases mAP and rank-1 by 3% and 7%, respectively, over ordinary triplet loss.

Effect of Color-Free Loss: Table VII illustrates the impact of \mathcal{L}_{cf} . The results indicate that incorporating this regularizing loss leads to an increase of 3% in mAP and 5% in rank-1 accuracy when it is used without using \mathcal{L}_{tri} or \mathcal{L}_{dual} . When the model is trained with \mathcal{L}_{cf} and \mathcal{L}_{tri} the performance significantly improves (19% in mAP and 12% in rank-1) and alongside with \mathcal{L}_{dual} mAP and rank-1 improved over 6% and 7% respectively. As the intermediate images are derived from visible images with color information removed and infrared style information incorporated, \mathcal{L}_{cf} plays a crucial role. This ensures that the extracted features from the visible modality are color-independent, aligning them more closely with the attributes of infrared.

Effect of Adversarial Loss: As a result of the adversarial loss, the generator must adapt the reconstruction process, enabling the stylistic transformation of images from V to I. Without using this objective, the generated images are similar to the gray-scale channel images and do help the model adjust for differences between I and V images. As you can see in the 4th row of Table. VII, the performance is close to the model trained based on grey intermediate images in [23].

Effect of Reconstruction Loss: To improve the quality of the reconstructed I images, supervised reconstruction loss is applied to the encoder and decoder parts. In the absence

Losses						R1 (%)	mAP (%)
\mathcal{L}_{id}	\mathcal{L}_{tri}	\mathcal{L}_{dual}	\mathcal{L}_{cf}	\mathcal{L}_{adv}	\mathcal{L}_{rec}		
✓	✓					49.41	41.55
✓		✓		✓	✓	60.64	56.15
✓				✓	✓	40.16	35.89
✓			✓	✓	✓	45.56	41.33
✓		✓	✓	✓	✓	60.67	58.35
✓		✓		✓		9.32	5.51
✓	✓		✓	✓	✓	62.76	60.02
✓	✓	✓	✓	✓	✓	64.80	61.92
✓		✓	✓	✓	✓	67.29	63.86

TABLE VII: Impact on the accuracy of using different losses in our AGPI² on SYSU-MM01 in single-shot.

of this loss, the generated images lack sufficient important information to support the feature module and make the performance drop dramatically. As the color-free loss pushes the V features to be the same as the features of such not meaningful intermediate images, the mAP accuracy of matching between V and I decreases (from 41.55% to 5.51%) as seen in the 5th row of Table VII.

3) *Generation Model*: To find the highest ReID performance with respect to the quality of the generated intermediate images, we tested different generative models by adapting V to I modality or vice versa. Since V images have more information to discriminate between individuals in comparison to I [68], the transferred V images play a better role in providing privileged information in the training process. In Table VIII, the model in [76] is deeper (more layers) and has a better capacity to stylize the V images from the I feature vectors. Also, we tested a large state-of-the-art model [81], which is used for multiple domain adaptation as our generator.

Model	V → I		I → V		# params
	R1	mAP	R1	mAP	
VQ-VAE2	65.75	61.96	61.81	58.34	20M
deep VQ-VAE2	67.29	63.86	63.45	60.09	28M
StarGAN	66.21	62.70	62.41	59.12	50M

TABLE VIII: Accuracy with different options for the generation of intermediate images.

4) ID-modality Discriminator vs General Discriminator:

The ID-modality discriminator aims to discriminate modalities based on ID-aware features for each modality to help the generative module transform such features from V and I to intermediate. To show the effectiveness of such a module on ReID performance, we compared its results with those of a general binary discriminator that uses whole features. Table IX demonstrates a 4% improvement in rank-1 accuracy and a 5% enhancement in mAP measurement by introducing our ID-modality discriminator. The discriminator effectively identifies modality-specific, ID-related information, enabling the generative model to transform visible images into infrared while preserving the individuals' discriminative attributes. Consequently, the transformed images provide intermediate privileged information to the feature backbone, leading to a

more robust reduction of modality discrepancies in the feature space. Additionally, Fig. 4 highlights the key regions of the input images that the discriminator focuses on when making its decisions, demonstrating the relevance of the identified features.

Model	V → I		I → V	
	R1	mAP	R1	mAP
Binary Discriminator	63.23	58.84	60.18	57.32
Our Discriminator	67.29	63.86	63.45	60.09

TABLE IX: Accuracy with different options for the discriminator and generation of intermediate images.

F. Inference Model Efficiency

Our AGPI² proposed method creates a privileged intermediate domain only in training to improve the ReID performance without adding extra computation overhead in the inference time. In Table X, we compared our model's size and time complexity in the inference time with the baseline [13], MPANet [48], MMN [24], DTRM [55], RPIG [20] and thGAN [12].

Model	# of Para. (M)	Flops (G)
baseline [13]	24.8	5.2
AGPI ²	24.8	5.2
MPANet [48]	30.1	6.7
MMN [24]	26.6	5.8
DTRM [55]	26.6	5.4
RPIG [20]	24.9	5.2
thGAN [12]	35	8.9

TABLE X: Number parameters and floating-point operations at inference time for AGPI² and state-of-the-art models.

G. Data Augmentation

Random erasing augmentation [82] was assessed with different settings for the feature embedding and the generation modules. Table XI shows the accuracy of our proposed approach using different scenarios [82]. Despite each augmentation configuration improving performance, the best results are achieved when the augmentation probability is gradually increased during training.

Configurations	p	s	R1 (%)	mAP (%)
AGPI ²	0	-	67.29	63.86
AGPI ² + Aug1	0.40	0.20	69.81	65.19
AGPI ² + Aug2	0.50	0.25	70.93	66.15
AGPI ² + Aug3	0.50	0.30	69.39	64.95
AGPI ² + Aug4	0.30 → 0.50	0.20 → 0.30	72.23	70.58

TABLE XI: Matching accuracy with random erasing augmentation. s is the size of the rectangle and p is the probability of erasing. Note: " $a \rightarrow b$ " means that training started from a and finished with b .

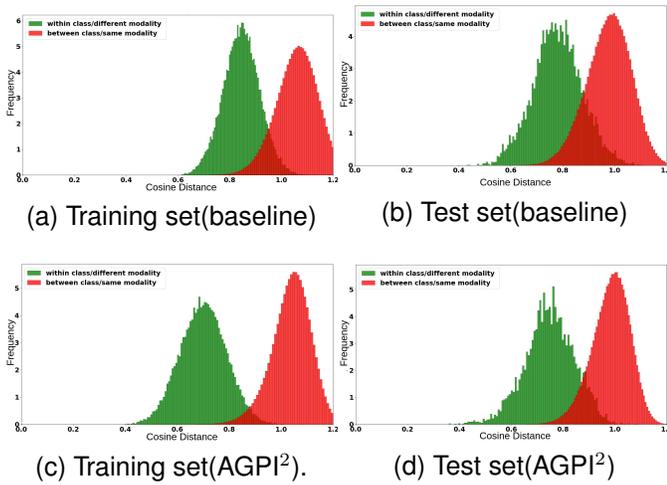


Fig. 5: Distribution of between- and within-class distances on the SYSU-MM01 dataset for (a,c) train and (b,d) test sets with the baseline model [13] and our AGPI² model.

H. Qualitative Evaluation

A visual representation of the cosine distance distributions for positive and negative cross-modality matching pairs is provided in Fig. 5 for both training and testing sets. The proposed AGPI² approach is compared to a strong baseline (AGW model in [13]). After training, such relationships should be reversed so that the two strategies work effectively to decouple I from V. Despite this, AGPI² is more efficient than the baseline on the test set when combined with the adaptively generated images. By leveraging privileged intermediate information during training, AGPI² improves the model’s ability to discriminate between V and I images at test time.

Fig. 6 shows the rank-10 for a given class of the SYSU-MM01 dataset given by the AGW [13] (the baseline) and our proposed strategy, along with the corresponding similarity CAMs [83]. The query CAM is produced from the similarity of the best match. Compared to the baseline method, most discriminative regions and attention between them are refined and more focused. CAMs seem to focus on shoulders, hips, and feet for the baseline method, while the hip’s attention almost fully disappears for our model. Interestingly, this behavior is consistent from one sample to another. Perhaps the hip focus at the baseline comes from the important color changes between the lower and upper bodies, which are probably discriminant for the visible part of the model. However, regarding the IR modality, this color marker is much less reliable and probably appears as a source of confusion for cross-modal ReID. This phenomenon tends to diminish thanks to the adaptive and informative intermediate privileged domain that is created through the training.

To demonstrate the effectiveness of our proposed ID-Modality discriminator, inspired by [41], [51], we compared its class activation map to that of a general binary discriminator, as shown in Fig. 4. The results indicate that the ID-Modality discriminator is more focused on the foreground, whereas the binary discriminator primarily detects the modality based on

the background.

V. CONCLUSION

In this paper, a new framework is introduced for adapting privileged intermediate information called AGPI². This method generates a virtual intermediate domain that bridges the gap between V and I images during training, thereby reducing the modality discrepancy problem. To create privileged intermediate information, an adversarial approach is also introduced for the generator, feature embedding, and ID-modality discriminator modules in the training process. AG outperforms state-of-the-art methods for V-I person ReID on the SYSU-MM01 and RegDB datasets while incurring no computational overhead during inference.

The AGPI² training approach can be integrated into other SOTA methods for ReID to improve performance and generality. Compared to other intermediate generation methods, we show AGPI² approach can provide more balance between V and I information for the model to diminish the modality-specific attributes in the representation feature space. In future research endeavours, we aim to extend the application of the AGPI² approach to multi-step and gradual intermediate privileged learning, particularly to address more substantial domain shift challenges. This will necessitate the AGPI² model’s ability to quantify the incorporation of I domain information during the transformation of V images into the intermediate space, facilitating a gradual adaptation process.

REFERENCES

- [1] S. He, W. Chen, K. Wang, H. Luo, F. Wang, W. Jiang, and H. Ding, “Region generation and assessment network for occluded person reid,” *IEEE Trans. Information Forensics and Security*, vol. 19, pp. 120–132, 2024. 1
- [2] D. Fu, D. Chen, J. Bao, H. Yang, L. Yuan, L. Zhang, H. Li, and D. Chen, “Unsupervised pre-training for person reid,” in *CVPR 2021*. 1
- [3] C. Sharma, S. R. Kapil, and D. Chapman, “Person reid with a locally aware transformer,” in *arXiv:2106.03720*, 2021. 1
- [4] V. Somers, C. De Vleeschouwer, and A. Alahi, “Body part-based representation learning for occluded person reid,” in *WACV 2023*. 1
- [5] A. Bhuiyan, Y. Liu, P. Siva, M. Javan, I. B. Ayed, and E. Granger, “Pose guided gated fusion for person reid,” in *WACV 2020*. 1
- [6] D. Mekhazni, A. Bhuiyan, G. Ekladiou, and E. Granger, “Unsupervised domain adaptation in the dissimilarity space for person reid,” in *ECCV 2020*. 1
- [7] Z. Lu, R. Lin, and H. Hu, “Modality and camera factors bi-disentanglement for nir-vis object reid,” *IEEE Trans. Information Forensics and Security*, vol. 18, pp. 1989–2004, 2023. 1, 4, 7
- [8] Y. Du, C. Lei, Z. Zhao, Y. Dong, and F. Su, “Video-based visible-infrared person reid with auxiliary samples,” *IEEE Trans. Information Forensics and Security*, vol. 19, pp. 1313–1325, 2024. 1
- [9] Y. Zhang, Y. Kang, S. Zhao, and J. Shen, “Dual-semantic consistency learning for visible-infrared person reid,” *IEEE Trans. Information Forensics and Security*, vol. 18, pp. 1554–1565, 2023. 1
- [10] M. Alehdaghi, P. Shamsolmoali, R. M. Cruz, and E. Granger, “Bidirectional multi-step domain generalization for visible-infrared person re-identification,” *arXiv preprint arXiv:2403.10782*, 2024. 1, 3
- [11] G. Wang, T. Zhang, J. Cheng, S. Liu, Y. Yang, and Z. Hou, “Rgb-infrared cross-modality person reid via joint pixel and feature alignment,” in *ICCV 2019*. 1, 2, 3, 7, 4
- [12] V. V. Kniaz, V. A. Knyaz, J. Hladuvka, W. G. Kropatsch, and V. Mizginov, “Thermalgan: Multimodal color-to-thermal image translation for person reid in multispectral dataset,” in *ECCV 2018*. 1, 2, 3, 10, 4
- [13] M. Ye, J. Shen, G. Lin, T. Xiang, L. Shao, and S. C. H. Hoi, “Deep learning for person reid: A survey and outlook,” *arXiv:2001.04193*, 2020. 1, 3, 5, 6, 7, 8, 10, 11

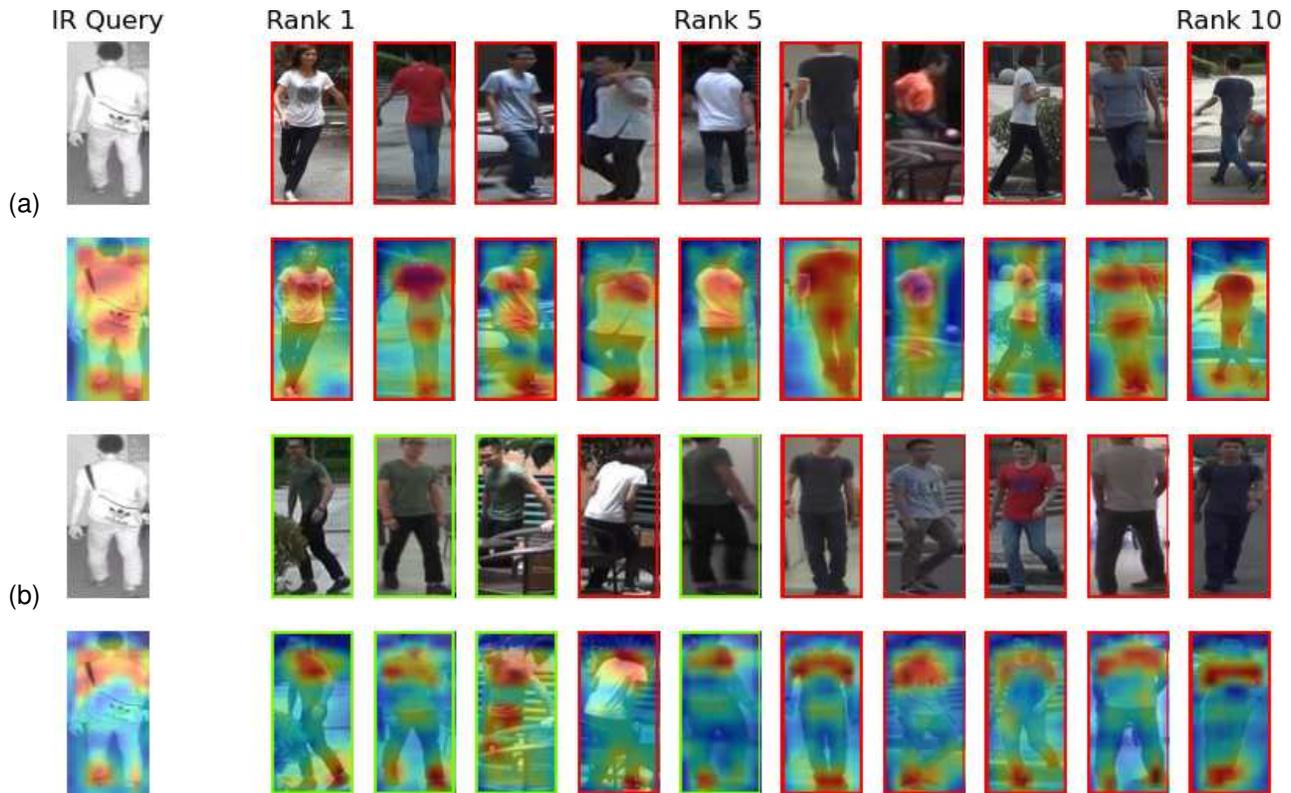


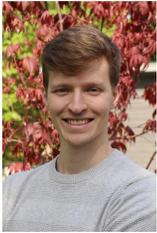
Fig. 6: Rank-10 visualizations of the (a) baseline (b) AGPI² strategies on SYSU-MM01 data, along with their corresponding similarity CAM [83]. Green or red boxes respectfully surround correct and incorrect matches. There are only four visible cameras, and only one image from one of them is in the gallery.

- [14] M. Ye, X. Lan, Z. Wang, and P. C. Yuen, "Bi-directional center-constrained top-ranking for visible thermal person reid," *IEEE Trans. Information Forensics and Security*, vol. 15, pp. 407–419, 2020. 1, 3
- [15] H. Liu and J. Cheng, "Enhancing the discriminative feature learning for visible-thermal cross-modality person reid," *CoRR*, vol. abs/1907.09659, 2019. [Online]. Available: <http://arxiv.org/abs/1907.09659> 1, 3
- [16] Y. Zhu, Z. Yang, L. Wang, S. Zhao, X. Hu, and D. Tao, "Hetero-center loss for cross-modality person reid," *Neurocomputing*, vol. 386, pp. 97–109, 2020. 1, 3
- [17] H. Park, S. Lee, J. Lee, and B. Ham, "Learning by aligning: Visible-infrared person reid using cross-modal correspondences," in *ICCV 2021*. 1, 3
- [18] A. Wu, W.-S. Zheng, H.-X. Yu, S. Gong, and J. Lai, "Rgb-infrared cross-modality person reid," in *ICCV 2017*. 1, 4
- [19] V. Vapnik and A. Vashist, "A new learning paradigm: Learning using privileged information," *Neural Networks*, vol. 22, no. 5-6, pp. 544–557, 2009. 1
- [20] M. Alehdaghi, A. Josi, R. M. Cruz, and E. Granger, "Visible-infrared person reid using privileged intermediate information," in *ECCVw 2022*. 1, 3, 7, 8, 9, 10, 2
- [21] D. Li, X. Wei, X. Hong, and Y. Gong, "Infrared-visible cross-modal person reid with an x modality," in *AAAI 2020*. 1, 4
- [22] Z. Wei, X. Yang, N. Wang, and X. Gao, "Syncretic modality collaborative learning for visible infrared person reid," in *ICCV 2021*. 1, 7
- [23] M. Ye, J. Shen, and L. Shao, "Visible-infrared person reid via homogeneous augmented tri-modal learning," *IEEE Trans. Information Forensics and Security*, vol. 16, pp. 728–739, 2020. 1, 3, 6, 7, 8, 9, 2
- [24] Y. Zhang, Y. Yan, Y. Lu, and H. Wang, "Towards a unified middle modality learning for visible-infrared person reid," in *ACM ICM 2021*. 1, 4, 7, 8, 10, 2
- [25] M. Alehdaghi, R. Bhattacharya, P. Shamsolmoali, R. M. O. Cruz, and E. Granger, "From cross-modal to mixed-modal visible-infrared re-identification," 2025. [Online]. Available: <https://arxiv.org/abs/2501.13307> 1
- [26] H. Yu, X. Cheng, W. Peng, W. Liu, and G. Zhao, "Modality unifying network for visible-infrared person re-identification," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 11 185–11 195. 3, 9
- [27] Y. Dai, J. Liu, Y. Sun, Z. Tong, C. Zhang, and L.-Y. Duan, "Idm: An intermediate domain module for domain adaptive person re-id," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 11 864–11 874. 3
- [28] A. Wu, W.-S. Zheng, S. Gong, and J. Lai, "Rgb-ir person reid by cross-modality similarity preservation," *International journal of computer vision*, vol. 128, no. 6, pp. 1765–1785, 2020. 3, 7
- [29] D. T. Nguyen, H. G. Hong, K. W. Kim, and K. R. Park, "Person recognition system based on a combination of body images from visible light and thermal cameras," *Sensors*, vol. 17, no. 3, p. 605, 2017. 3, 7
- [30] Y. Chen, L. Wan, Z. Li, Q. Jing, and Z. Sun, "Neural feature search for rgb-infrared person reid," in *CVPR 2021*. 3
- [31] P. Dai, R. Ji, H. Wang, Q. Wu, and Y. Huang, "Cross-modality person reid with generative adversarial training," in *IJCAI 2018*. 3
- [32] C. Fu, Y. Hu, X. Wu, H. Shi, T. Mei, and R. He, "Cm-nas: Cross-modality neural architecture search for visible-infrared person reid," in *ICCV 2021*. 3
- [33] Y. Hao, J. Li, N. Wang, and X. Gao, "Modality adversarial neural network for visible-thermal person reid," *Pattern Recognition*, vol. 107, p. 107533, 2020. 3
- [34] X. Xu, S. Wu, S. Liu, and G. Xiao, "Cross-modal based person reid via channel exchange and adversarial learning," in *NeurIPS 2021*. 3
- [35] M. Ye, X. Lan, J. Li, and P. Yuen, "Hierarchical discriminative learning for visible thermal person reid," in *AAAI 2018*. 3
- [36] Z. Wang, Z. Wang, Y. Zheng, Y.-Y. Chuang, and S. Satoh, "Learning to reduce dual-level discrepancy for infrared-visible person reid," in *CVPR 2019*. 3, 7, 2
- [37] S. Choi, S. Lee, Y. Kim, T. Kim, and C. Kim, "Hi-cmd: hierarchical cross-modality disentanglement for visible-infrared person reid," in *CVPR 2020*. 3, 7, 2, 4
- [38] G.-A. Wang, T. Zhang, Y. Yang, J. Cheng, J. Chang, X. Liang, and Z.-G. Hou, "Cross-modality paired-images generation for rgb-infrared person reid," in *AAAI 2020*. 3, 7

- [39] Y. Yang, T. Zhang, J. Cheng, Z. Hou, P. Tiwari, H. M. Pandey *et al.*, “Cross-modality paired-images generation and augmentation for rgb-infrared person reid,” *Neural Networks*, vol. 128, pp. 294–304, 2020. **3**
- [40] A. Wu, W.-S. Zheng, H.-X. Yu, S. Gong, and J. Lai, “Rgb-infrared cross-modality person reid,” in *ICCV 2021*. **3**
- [41] L. Wu, D. Liu, W. Zhang, D. Chen, Z. Ge, F. Boussaid, M. Bennamoun, and J. Shen, “Pseudo-pair based self-similarity learning for unsupervised person re-identification,” *IEEE Transactions on Image Processing*, vol. 31, pp. 4803–4816, 2022. **3, 11**
- [42] M. Ye, W. Ruan, B. Du, and M. Z. Shou, “Channel augmented joint learning for visible-infrared recognition,” in *ICCV 2021*. **3, 7**
- [43] J. Jiang, K. Jin, M. Qi, Q. Wang, J. Wu, and C. Chen, “A cross-modal multi-granularity attention network for rgb-ir person reid,” *Neurocomputing*, 2020. **3**
- [44] K. Chen, Z. Pan, J. Wang, S. Jiao, Z. Zeng, and Z. Miao, “Joint feature learning network for visible-infrared person reid,” in *PRCV 2020*. **3**
- [45] M. Ye, J. Shen, D. J. Crandall, L. Shao, and J. Luo, “Dynamic dual-attentive aggregation learning for visible-infrared person reid,” in *ECCV 2020*. **3, 6, 7, 8**
- [46] Y. Lu, Y. Wu, B. Liu, T. Zhang, B. Li, Q. Chu, and N. Yu, “Cross-modality person reid with shared-specific feature transfer,” in *CVPR 2020*. **3, 7**
- [47] Z. Huang, J. Liu, L. Li, K. Zheng, and Z.-J. Zha, “Modality-adaptive mixup and invariant decomposition for rgb-infrared person reid,” in *AAAI 2022*. **3**
- [48] Q. Wu, P. Dai, J. Chen, C.-W. Lin, Y. Wu, F. Huang, B. Zhong, and R. Ji, “Discover cross-modality nuances for visible-infrared person reid,” in *CVPR 2021*. **3, 7, 8, 10**
- [49] X. Fang, Y. Yang, and Y. Fu, “Visible-infrared person reid via semantic alignment and affinity inference,” in *ICCV 2023*. **3**
- [50] M. Kim, S. Kim, J. Park, S. Park, and K. Sohn, “Partmix: Regularization strategy to learn part discovery for visible-infrared person reid,” in *CVPR 2023*. **3**
- [51] L. Y. Wu, L. Liu, Y. Wang, Z. Zhang, F. Boussaid, M. Bennamoun, and X. Xie, “Learning resolution-adaptive representations for cross-resolution person re-identification,” *IEEE Transactions on Image Processing*, 2023. **3, 11**
- [52] J. Wu, H. Liu, Y. Su, W. Shi, and H. Tang, “Learning concordant attention via target-aware alignment for visible-infrared person reid,” in *ICCV 2023*. **3**
- [53] X. Fan, H. Luo, C. Zhang, and W. Jiang, “Cross-spectrum dual-subspace pairing for rgb-infrared cross-modality person reid,” *ArXiv*, vol. abs/2003.00213, 2020. **3**
- [54] H. Liu, S. Ma, D. Xia, and S. Li, “Sfanet: A spectrum-aware feature augmentation network for visible-infrared person reidentification,” *IEEE Trans. Neural Networks and Learning Systems*, 2021. **3**
- [55] M. Ye, C. Chen, J. Shen, and L. Shao, “Dynamic tri-level relation mining with attentive graph for visible infrared reid,” *IEEE Trans. Information Forensics and Security*, vol. 17, pp. 386–398, 2022. **3, 7, 10**
- [56] D. Pechyony and V. Vapnik, “On the theory of learning with privileged information,” *NeurIPS 2010*. **4, 5, 6**
- [57] J. Hoffman, S. Gupta, and T. Darrell, “Learning with side information through modality hallucination,” in *CVPR 2016*. **4**
- [58] C. Choi, S. Kim, and K. Ramani, “Learning hand articulations by hallucinating heat distribution,” in *ICCV 2017*. **4**
- [59] M. Kampffmeyer, A.-B. Salberg, and R. Jenssen, “Urban land cover classification with missing data modalities using deep convolutional neural networks,” *IEEE J. Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 11, no. 6, pp. 1758–1768, 2018. **4**
- [60] S. Kumar, B. Banerjee, and S. Chaudhuri, “Improved landcover classification using online spectral data hallucination,” *Neurocomputing*, vol. 439, pp. 316–326, 2021. **4**
- [61] J. Lezama, Q. Qiu, and G. Sapiro, “Not afraid of the dark: Nir-vis face recognition via cross-spectral hallucination and low-rank embedding,” in *CVPR 2017*. **4**
- [62] M. R. U. Saputra, P. P. de Gusmao, C. X. Lu, Y. Almalioglu, S. Rosa, C. Chen, J. Wahlström, W. Wang, A. Markham, and N. Trigoni, “Deeptio: A deep thermal-inertial odometry with visual hallucination,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1672–1679, 2020. **4**
- [63] S. Pande, A. Banerjee, S. Kumar, B. Banerjee, and S. Chaudhuri, “An adversarial approach to discriminative modality distillation for remote sensing image classification,” in *ICCVw 2019*. **4**
- [64] N. Crasto, P. Weinzaepfel, K. Alahari, and C. Schmid, “Mars: Motion-augmented rgb stream for action recognition,” in *CVPR 2019*. **4**
- [65] S. Chopra, R. Hadsell, and Y. LeCun, “Learning a similarity metric discriminatively, with application to face verification,” in *CVPR 2005*. **4**
- [66] F. Schroff, D. Kalenichenko, and J. Philbin, “Facenet: A unified embedding for face recognition and clustering,” in *CVPR 2015*. **4**
- [67] J. Lambert, O. Sener, and S. Savarese, “Deep learning under privileged information using heteroscedastic dropout,” in *CVPR 2018*. **4**
- [68] M. Jia, Y. Zhai, S. Lu, S. Ma, and J. Zhang, “A similarity inference metric for rgb-infrared cross-modality person reid,” *arXiv:2007.01504*, 2020. **6, 7, 10**
- [69] X. Wang, R. Girshick, A. Gupta, and K. He, “Non-local neural networks,” in *CVPR 2018*. **6**
- [70] Z. Zhang, S. Jiang, C. Huang, Y. Li, and R. Y. Da Xu, “Rgb-ir cross-modality person reid based on teacher-student gan model,” *Pattern Recognition Letters*, vol. 150, pp. 155–161, 2021. **7**
- [71] J. Sun, Y. Li, H. Chen, Y. Peng, and J. Zhu, “Visible-infrared person reid model based on feature consistency and modal indistinguishability,” *Machine Vision and Applications*, vol. 34, no. 1, 2023. **7**
- [72] X. Zeng, J. Long, S. Tian, and G. Xiao, “Random area pixel variation and random area transform for visible-infrared cross-modal pedestrian reid,” *Expert Systems with Applications*, vol. 215, p. 119307, 2023. **7**
- [73] “G2da: Geometry-guided dual-alignment learning for rgb-infrared person reid,” *Pattern Recognition*, vol. 135, p. 109150, 2023. **7**
- [74] Y. Zhang and H. Wang, “Diverse embedding expansion network and low-light cross-modality benchmark for visible-infrared person re-identification,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 2153–2162. **7, 3, 4**
- [75] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *CVPR 2016*. **7**
- [76] A. Van Den Oord, O. Vinyals *et al.*, “Neural discrete representation learning,” *NeurIPS 2017*. **7, 10**
- [77] J. Feng, A. Wu, and W.-S. Zheng, “Shape-erased feature learning for visible-infrared person reid,” in *CVPR 2023*. **8**
- [78] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, “A kernel two-sample test,” *J. Machine Learning Research*, vol. 13, no. 1, pp. 723–773, 2012. **8**
- [79] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *CVPR 2009*. **8**
- [80] J. Gildenblat and contributors, “Pytorch library for cam methods,” <https://github.com/jacobgil/pytorch-grad-cam>, 2021. **9**
- [81] Y. Choi, Y. Uh, J. Yoo, and J.-W. Ha, “Stargan v2: Diverse image synthesis for multiple domains,” in *CVPR 2020*. **10**
- [82] Z. Zhong, L. Zheng, G. Kang, S. Li, and Y. Yang, “Random erasing data augmentation,” in *AAAI 2020*. **10**
- [83] A. Stylianou, R. Souvenir, and R. Pless, “Visualizing deep similarity networks,” in *WACV 2019*. **11, 12**
- [84] M. Boudiaf, J. Rony, I. M. Ziko, E. Granger, M. Pedersoli, P. Piantanida, and I. B. Ayed, “A unifying mutual information view of metric learning: cross-entropy vs. pairwise losses,” in *European conference on computer vision*, 2020, pp. 548–564. **1**
- [85] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in *ICCV*, 2017. **1, 2, 3, 4**
- [86] L. McInnes, J. Healy, and J. Melville, “Umap: Uniform manifold approximation and projection for dimension reduction,” *arXiv:1802.03426*, 2018. **3**



Mahdi Alehdaghi received an M.S. degree in artificial intelligence in 2015 from the Ferdowsi University of Mashhad, Mashhad-Iran. He is currently pursuing a Ph.D. degree in computer engineering from the École de Technologie Supérieure (ÉTS), Montréal, Quebec. His current research interests include computer vision, multimodal representations, and gradual and multi-step domain generalization.



Arthur Josi received an M.Sc degree in artificial intelligence in 2023 from the 'École de technologie supérieure (ÉTS)' Montréal, Québec. He is currently pursuing a PhD in computer engineering from the ÉTS. His current research interests are in computer vision, multimodal fusion, 3D facial animation, and deep generative models.



Pourya Shamsolmoali (Senior Member, IEEE) received the Ph.D. degree in computer science from Shanghai Jiao Tong University. He has been a Visiting Researcher with CMCC-Italy, INRIA-France, and ETS-Canada. He is currently a lecturer at the University of York with expertise in image processing, computer vision, and machine learning.



Rafael M. O. Cruz (Member, IEEE) received a Ph.D. in engineering in 2016 from the École de Technologie Supérieure (ÉTS), Montréal-Quebec, where he is currently an Associate Professor. His current research interests include ensemble learning, data complexity, dynamic ensemble models, meta-learning, and biometrics.



Eric Granger (Member, IEEE) received a Ph.D. degree in electrical engineering from École Polytechnique de Montréal in 2001. He was a Defense Scientist with DRDC Ottawa, from 1999 to 2001, and in Research and Development with Mitel Networks from 2001 to 2004. He joined the Department of Systems Engineering, ETS Montréal, Canada, in 2004, where he is currently a Full Professor and the Director of the LIVIA. He is the FRQS Co-Chair in AI and Health, and the ETS Industrial Research Co-Chair on embedded neural networks for intelligent

connected buildings (Distech Controls Inc.). His research interests include pattern recognition, machine learning, information fusion, and computer vision, with applications in affective computing, biometrics, face recognition, medical image analysis, and video surveillance.

APPENDIX A PROOFS

1. Maximizing $MI(Z; Y)$

Proposition 1: Let Z and Y be random variables with domains \mathcal{Z} and \mathcal{Y} , respectively. Minimizing the conditional cross-entropy loss of predicted label \hat{Y} , denoted by $\mathcal{H}(Y; \hat{Y}|Z)$, is equivalent to maximizing the $MI(Z; Y)$

Proof 1: Let us define the MI as entropy,

$$MI(Z; Y) = \underbrace{\mathcal{H}(Y)}_{\delta} - \underbrace{\mathcal{H}(Y|Z)}_{\xi} \quad (\text{A.1})$$

Since the domain \mathcal{Y} does not change, the entropy of the identity δ term is a constant and can therefore be ignored. Maximizing $MI(Z; Y)$ can only be achieved by minimizing the ξ term. By expanding its relation to the cross-entropy [84]:

$$\mathcal{H}(Y; \hat{Y}|Z) = \mathcal{H}(Y|Z) + \underbrace{\mathcal{D}_{KL}(Y||\hat{Y}|Z)}_{\geq 0}, \quad (\text{A.2})$$

where:

$$\mathcal{H}(Y|Z) \leq \mathcal{H}(Y; \hat{Y}|Z). \quad (\text{A.3})$$

So, $\mathcal{H}(Y|Z)$ is upper-bounded by our cross-entropy loss, and minimizing such loss results in minimizing the ξ term.

Through the maximization ID-Aware part of Eq. 3, training can naturally be decoupled in 2 steps. First, weights of the generative and feature backbone modules are fixed, and only the classifier parameters (i.e., the weight of the fully connected layer at the feature backbone) are minimized w.r.t. Eq. A.2. Through this step, $\mathcal{D}_{KL}(Y||\hat{Y}|Z)$ is minimized by adjusting \hat{Y} while the $\mathcal{H}(Y|Z)$ does not change. In the second step, the prototype module's weights are minimized w.r.t. $\mathcal{H}(Y|Z)$, while the classifier parameters are fixed.

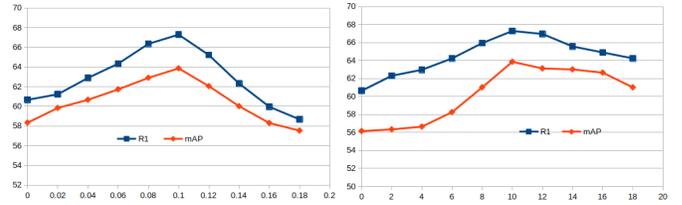
APPENDIX B ADDITIONAL RESULTS

1. Hyperparameter Values

This subsection provides an analysis of the impact on ReID accuracy of the λ_a and λ_{cf} hyperparameters. λ_a balances how much the model preserves the content of input images and the information of the style image, whereas λ_{cf} pushes the similarity of the intermediate feature on V extracted features. Fig. B.1 a) shows that accuracy improves when λ_a increases until it reaches a value of 0.1. Then, the performance begins to decline. Similar results can be observed when λ_{cf} varies from 1 to 10 (Fig. B.1 b)). The performance of our proposed AGPI² strategies varies depending on these hyperparameter values, but most results exceed the state-of-the-art methods.

2. Comparison to CycleGAN

we conducted a new experiment to show the comparison between using intermediate and directly generated modalities in the training and inference phases separately. In Table. B.1, in the first two rows, the results for AGPI² are reported in which the generated images are used as an intermediate. The \mathcal{L}_{dual} has been applied in the training process. In comparison, in the two last rows, the generated images are used as



(a) Accuracy vs. λ_{adv} value. (b) Accuracy vs. λ_{cf} value.

Fig. B.1: Impact on accuracy of λ_{adv} and λ_{cf} hyperparameters on the SYSU-MM01 dataset under the single-shot setting.

Method	Use		Performance		Complexity	
	Train	Test	R1(%)	mAP(%)	# of Para.	Flops
Using generative images as intermediate in the training						
AGPI ²	✓	✗	72.23	70.58	24.8M	5.2G
AGPI ²	✓	✓	65.76	62.23	52.8M	8.3G
Using generative images as other modality in the training						
CycleGAN	✓	✗	42.05	41.17	24.8M	5.2G
CycleGAN	✓	✓	26.91	28.75	42.5M	7.1G

TABLE B.1: Impact of using generated images as intermediate versus using as visible in training or testing for SYSU dataset. AGPI² method uses them as intermediate in \mathcal{L}_{dual} and CycleGAN [85] uses as visible modality and applies the \mathcal{L}_{tri} .

other modality (Visible images are replaced with generated images in training) and the \mathcal{L}_{tri} is applied between I and generated images. In the second and fourth rows, the generated images are used as the gallery instead of the visible images. AGPI², which generates and uses images as an intermediate, performs better than CycleGAN [85] using generated images directly as another modality. We can conclude that although the transformed images try to lose style information from the source and achieve them from the target domain, they are likely to have some content information from the target and style from the source, so the provided privileged information by them is not enough for the model to compensate the other modality in the training.

Additionally, generating the transformed images in testing makes the model trigger the GAN module and increases the model complexity. Also, such generated images may lose some information compared to the original images and degrade the matching performance. These images are shown in Fig. B.2.

3. Feature Visualization

To demonstrate that the intermediate images in AGPI² training approach bridge between visible and infrared extracted features, we visualized the feature distributions of the intermediate, visible, and infrared representations during the training process. As shown in Fig. B.3, the intermediate features are positioned between the visible and infrared features. The visible features gradually align with the intermediate feature space when we apply \mathcal{L}_{cf} and the first part of \mathcal{L}_{dual} . As the model optimizes the second part of \mathcal{L}_{dual} , the infrared features progressively discard modality-specific information and converge toward the intermediate feature space. By the end of the training, the visible and infrared features had converged



Fig. B.2: Examples of images generated with CycleGAN [85] for SYSU-MM01 dataset. Column (a) contains the real V, and column (b) is the translated visible image to infrared. Columns (c) and (d) are the real I images and translated versions of them to the V domain.

to a small region of the feature space, indicating that they had lost their modality-discriminative characteristics.

4. MMD Distance on extracted learned features:

The objective of the intermediate feature is to serve as a bridge between visible and infrared information by providing privileged identity-discriminative information while eliminating modality-specific style information. To demonstrate the reduction in the domain gap achieved through the use of intermediate features, we applied the kernel Maximum Mean Discrepancy (MMD) function to the extracted features from visible, infrared, and intermediate images. To measure MMD between modalities, we compute the center feature vector for each person within each modality and apply the MMD kernel to these vectors. Since, during the inference phase, infrared images are used as queries to the model in order to identify the matched person from the gallery, the distance between infrared (I) and intermediate (Z) features is crucial. As shown in Table B.2, the use of intermediate features significantly reduces this distance compared to the baseline.

5. Comparing to Image-based Intermediate method in VI-ReID

We also conducted a new experiment showing the qualitative and quantitative comparison between CycleGAN [85], ThermalGAN [12], D²RL [36], HI-CMD [37] and MMN [24] and our AGPI². We used the official code published by the authors of these methods and trained them with default hyperparameters.

Modality	Train			Test		
	I-Z	V-Z	I-V (\downarrow)	I-Z	V-Z	I-V (\downarrow)
baseline	-	-	0.12	-	-	0.53
Grayscale (HAT [23])	0.07	0.03	0.08	0.46	0.09	0.45
Random(RPIG [20])	0.05	0.04	0.06	0.41	0.11	0.38
AGPI ² (ours)	0.03	0.05	0.05	0.38	0.16	0.33

TABLE B.2: MMD between V, I, and different intermediate learned features for 100 and 50 different persons from train and test sets, respectively, on the SYSU-MM01. Z images are the intermediate images reconstructed by \mathcal{G} .

Method	Use		Performance		Complexity	
	Train	Test	RI(%)	mAP(%)	#Para.	Flops
CycleGAN [85]	✓	✗	42.05	41.17	24.8M	5.2G
CycleGAN	✓	✓	26.91	28.75	42.5M	7.1G
thGAN [12]	✓	✗	-	-	-	-
thGAN	✓	✓	17.77	18.13	70.0M	13.7G
AlignGAN [11]	✓	✗	24.81	22.05	23.5M	5.2G
AlignGAN	✓	✓	42.4	40.7	50.4M	9.6G
HiCMD [37]	✓	✓	34.94	35.94	40.1M	7.3G
HiCMD	✓	✓	15.57	17.38	100.6M	13.5G
MMN [24]	✓	✗	65.14	61.57	72.9M	6.1G
MMN(Only Z)	✓	✓	62.96	58.34	35.4M	6.2G
MMN (Z,I,V)	✓	✓	70.1	66.2	72.9M	12.1G
AGPI ²	✓	✗	72.23	70.58	24.8M	5.2G
AGPI ²	✓	✓	65.76	62.23	52.8M	8.3G

TABLE B.3: Impact of using generated images as intermediate versus using as visible in training or testing for SYSU dataset. AGPI² method uses them as intermediate in \mathcal{L}_{dual} and CycleGAN [85] uses as visible modality and applies the \mathcal{L}_{tri} .

Table B.3 presents a comparison of open-sourced generative image-based methods for Visible-Infrared Person Re-Identification (VI-ReID) on the SYSU-MM01 dataset. For each method, we report the performance in terms of Rank-1 accuracy and mAP, alongside the computational complexity measured by the number of parameters and FLOPs. The table highlights the differences between using generated images as an intermediate representation during training versus as a matching modality during testing.

The MMN [24] approach utilizes generated intermediate images during both training and testing. Specifically, MMN extracts features from visible, infrared, and generated (Z) images, combining them for improved matching. While this strategy boosts performance, as evidenced by the 70.1% Rank-1 accuracy and 66.2% mAP in the (Z, I, V) setup, it significantly increases computational overhead, with a FLOPs count of 12.1G—almost double that of single-modality models.

In contrast, our proposed AGPI² method treats generated images as privileged information during training, incorporating them into the dual loss \mathcal{L}_{dual} , without altering the backbone for testing. This leads to significantly lower computational overhead, with only 8.3G FLOPs in the test phase, while maintaining competitive performance. Notably, when the generated images (Z) are used as a query instead of infrared images (I), AGPI² achieves a Rank-1 accuracy of 65.76%, outperforming MMN (62.96%), which demonstrates the informativeness and bridging capability of the generated images.

Furthermore, AGPI² not only delivers superior matching

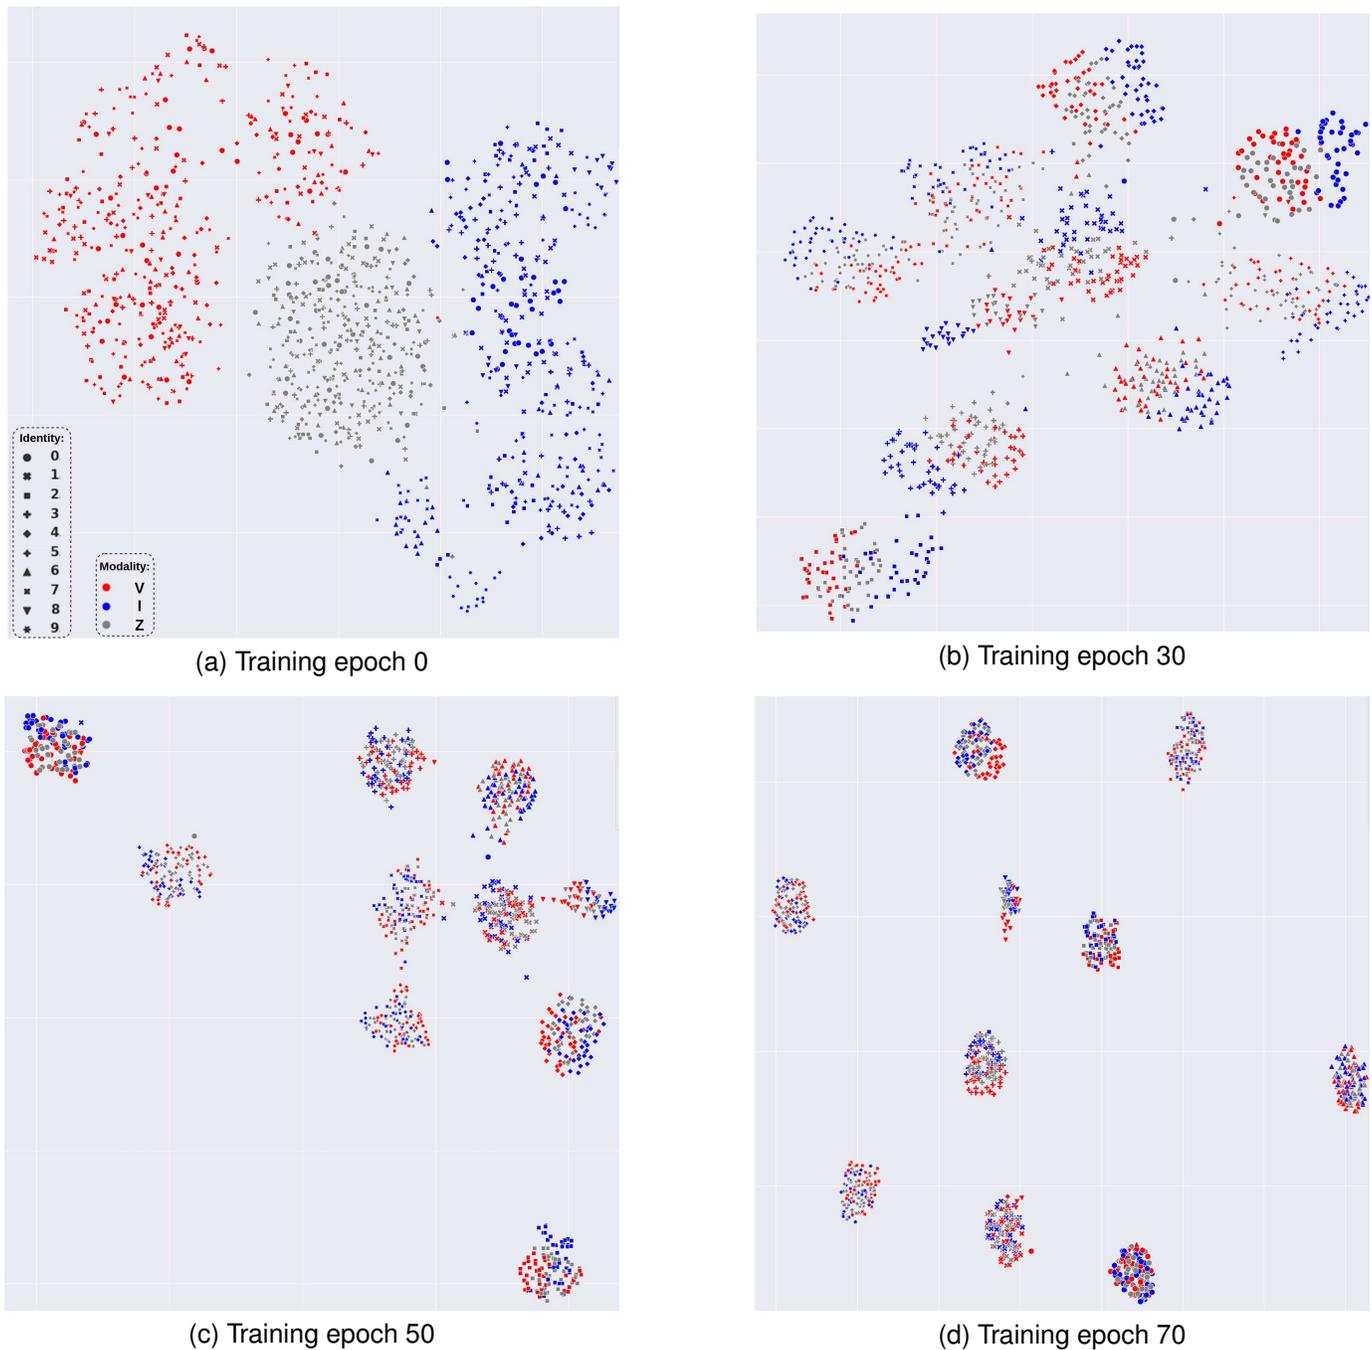


Fig. B.3: UMAP [86] visual representation of ten randomly selected identities from the SYSU-MM01 dataset through training process. Identities are shown with different shapes, and modalities are shown with different colors (red and blue for visible and infrared and intermediate features are shown in gray color).

performance but also minimizes computational costs compared to methods like CycleGAN [85] and thermalGAN [12], which require substantial resources due to their use of generated images during both training and testing. Overall, AGPI² balances effectiveness and efficiency, making it a compelling solution for VI-ReID tasks with constrained computational budgets.

Also, we generated some samples from the SYSU-MM01 dataset in Figure B.4 for each method. Our AGPI² produces more realistic images with more details compared to others.

6. Incorporating to Extra SOTA Models

Moreover, to compare with the requested method (PartMix [50] and DEEN [74]), we apply our AGPI² approach to them and report the performance and efficacy. The PartMIX method [50] utilizes part-based information by introducing an additional sub-module for part detection. To ensure a fair comparison, we integrated our AGPI² training approach into the SAAI method [49], a publicly available part-based VI-ReID method, and retrained it from scratch. Similarly,

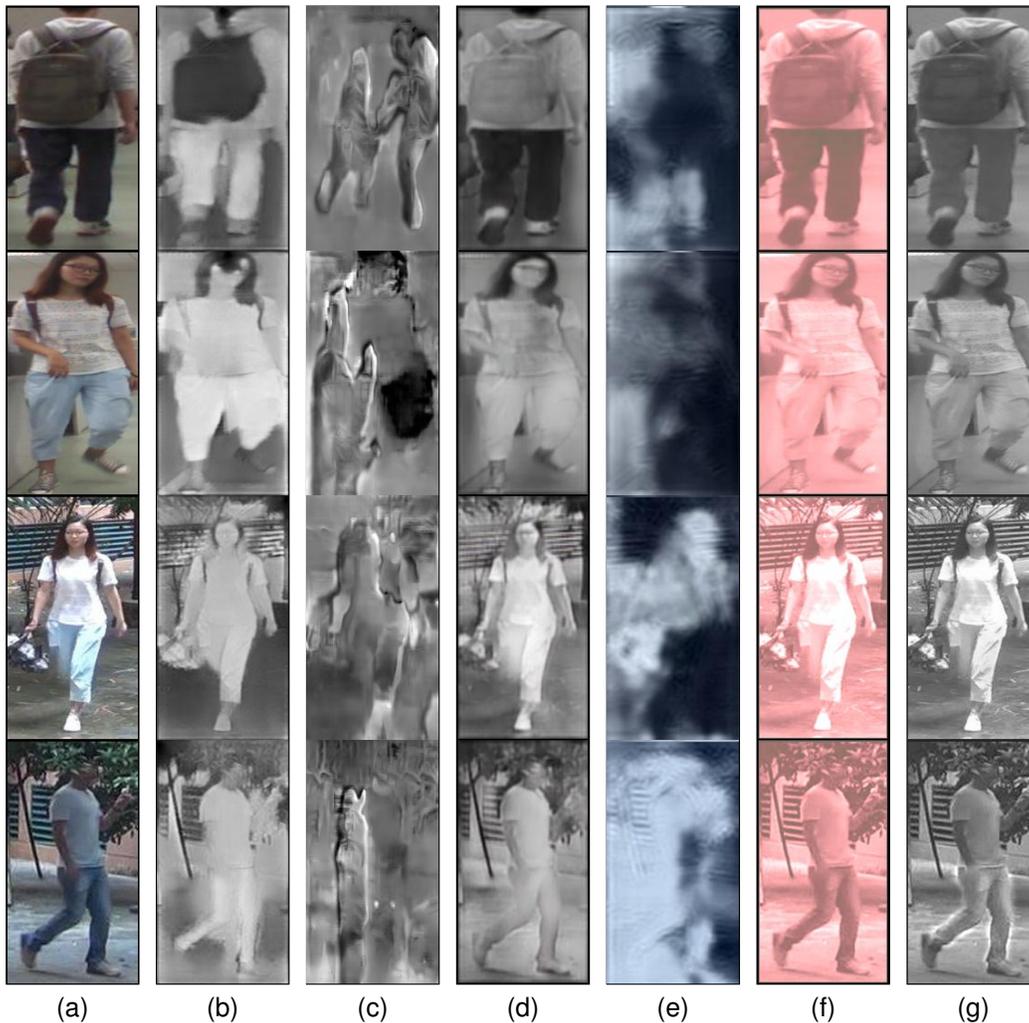


Fig. B.4: Examples of 4 intermediate images generated from (a) Visible by (b)CycleGAN [85],(c) ThermalGAN [12], (d) AlignGAN [11], (e) HiCMD [37], (f) MMN [24] and (g) our AGPI² on the SYSU-MM01 dataset.

AGPI² was added to the DEEN method [74] to evaluate its effectiveness on a non-part-based baseline. The results of these integrations on the SYSU-MM01 and RegDB datasets are shown in Table B.4.

For the part-based baseline, integrating AGPI² with part information leads to a substantial improvement in performance. Specifically, the combined AGPI² + Part approach achieves a notable increase in mAP for SYSU-MM01 (All Search) from 74.62% to 78.56% while maintaining competitive Rank-1 accuracy. On RegDB, the combined approach further improves cross-modal performance, with a significant boost in mAP for the I→V setting (from 82.52% to 91.15%). These results highlight the complementary nature of AGPI² when combined with part-based strategies, effectively enhancing both identity discrimination and modality bridging.

For the DEEN baseline, AGPI² integration results in consistent improvements across all settings. For example, in SYSU-MM01 (All Search), the DEEN + AGPI² combination improves mAP from 71.8% to 72.7%, and on RegDB (V→I), mAP increases from 85.8% to 86.41%. These gains

demonstrate the adaptability of AGPI² in improving feature representations across different baseline architectures. Notably, AGPI² achieves these enhancements with significantly lower complexity compared to other methods, as evidenced by its reduced parameter count and FLOPs.

These improvements demonstrate that AGPI² acts as a powerful enhancement tool when integrated with other models, further boosting their accuracy. Moreover, this integration does not introduce additional computational overhead during inference, making it a highly efficient solution.

Method	SYSU-MM01				RegDB				Complexity	
	All Search		Indoor Search		V \rightarrow I		I \rightarrow V		#Params	FLOPs
	R1	mAP	R1	mAP	R1	mAP	R1	mAP		
PartMix	77.78	74.62	81.52	84.34	84.93	82.52	85.66	82.27	52.5M	7.4G
AGPI ² (ours)	72.23	70.58	83.45	84.25	89.03	83.89	87.91	83.04	24.8M	5.2G
AGPI ² + Part	77.35	78.56	84.60	88.91	91.22	91.15	91.3	92.5	52.5M	7.4G
DEEN	74.7	71.8	80.3	83.3	91.1	85.8	89.5	83.4	89.0M	39.8G
AGPI ² (ours)	72.23	70.58	83.45	84.25	89.03	83.89	87.91	83.04	24.8M	5.2G
DEEN + AGPI ²	75.36	72.7	83.19	84.91	92.5	86.41	90.38	84.7	89.0M	39.8G

TABLE B.4: Accuracy of integrating the proposed AGPI² with different baselines on the SYSU-MM01 (single-shot setting) and RegDB datasets.