



Universiteit
Leiden
The Netherlands

Dual Gaussian-based variational subspace disentanglement for visible-infrared person re-identification

Pu, N.; Chen, W.; Liu, Y.; Bakker, E.M.; Lew, M.S.K.; Atrey, P.K.; Li, Z.

Citation

Pu, N., Chen, W., Liu, Y., Bakker, E. M., & Lew, M. S. K. (2020). Dual Gaussian-based variational subspace disentanglement for visible-infrared person re-identification. *Mm '20: Proceedings Of The 28Th Acm International Conference On Multimedia*, 2149-2158. doi:10.1145/3394171.3413673

Version: Publisher's Version

License: [Licensed under Article 25fa Copyright Act/Law \(Amendment Taverne\)](#)

Downloaded from: <https://hdl.handle.net/1887/3418406>

Note: To cite this publication please use the final published version (if applicable).

Dual Gaussian-based Variational Subspace Disentanglement for Visible-Infrared Person Re-Identification

Nan Pu
n.pu@liacs.leidenuniv.nl
LIACS Media Lab, Leiden University

Wei Chen
w.chen@liacs.leidenuniv.nl
LIACS Media Lab, Leiden University

Yu Liu
yu.liu@esat.kuleuven.be
ESAT-PSI, KU Leuven University

Erwin M. Bakker
erwin@liacs.leidenuniv.nl
LIACS Media Lab, Leiden University

Michael S. Lew
m.s.k.lew@liacs.leidenuniv.nl
LIACS Media Lab, Leiden University

ABSTRACT

Visible-infrared person re-identification (VI-ReID) is a challenging and essential task in night-time intelligent surveillance systems. Except for the intra-modality variance that RGB-RGB person re-identification mainly overcomes, VI-ReID suffers from additional inter-modality variance caused by the inherent heterogeneous gap. To solve the problem, we present a carefully designed *dual Gaussian-based variational auto-encoder* (DG-VAE), which disentangles an identity-discriminable and an identity-ambiguous cross-modality feature subspace, following a mixture-of-Gaussians (MoG) prior and a standard Gaussian distribution prior, respectively. Disentangling cross-modality identity-discriminable features leads to more robust retrieval for VI-ReID. To achieve efficient optimization like conventional VAE, we theoretically derive two variational inference terms for the MoG prior under the supervised setting, which not only restricts the identity-discriminable subspace so that the model explicitly handles the cross-modality intra-identity variance, but also enables the MoG distribution to avoid posterior collapse. Furthermore, we propose a *triplet swap reconstruction* (TSR) strategy to promote the above disentangling process. Extensive experiments demonstrate that our method outperforms state-of-the-art methods on two VI-ReID datasets. Codes will be available at <https://github.com/TPCD/DG-VAE>.

CCS CONCEPTS

• Information systems → Information retrieval.

KEYWORDS

cross-modality person re-identification; variational auto-encoder; disentangled representation; Gaussian mixture model

ACM Reference Format:

Nan Pu, Wei Chen, Yu Liu, Erwin M. Bakker, and Michael S. Lew. 2020. Dual Gaussian-based Variational Subspace Disentanglement for Visible-Infrared Person Re-Identification. In *28th ACM International Conference on Multimedia (MM '20)*, October 12–16, 2020, Seattle, WA, USA. ACM, New York, NY, USA, 20 pages. <https://doi.org/10.1145/3394171.3413673>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

MM '20, October 12–16, 2020, Seattle, WA, USA.

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7988-5/20/10...\$15.00

<https://doi.org/10.1145/3394171.3413673>

1 INTRODUCTION

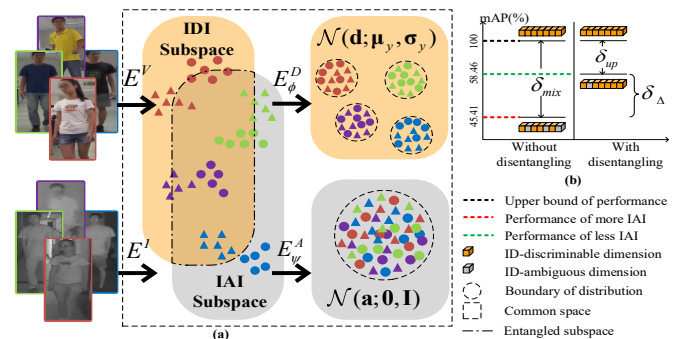


Figure 1: (a) Conceptual illustration of the dual Gaussian-based variational subspace disentanglement. Our DG-VAE factorizes identity-discriminable information (IDI) and identity-ambiguous information (IAI) from the entangled common space of visible-infrared features. The IDI follows the mixture-of-Gaussians prior and the IAI follows the standard Gaussian distribution prior. (b) shows the effectiveness of disentangling methods. After excluding the IAI from the learned representation, DG-VAE approximates the degraded performance (δ_{mix}) to the IDI-based performance (δ_{up}).

Person re-identification (Re-ID), which aims at associating the same pedestrian images across disjoint camera views, has received ever-increasing attention from the computer vision community [22, 40, 50]. Most recent efforts [2–4, 13, 20, 24, 43, 44, 49, 51] are focused on single-modality image retrieval, e.g., RGB-RGB image matching, which depends on good visible light conditions to catch the appearance of pedestrian. However, visible images might have inferior quality due to inadequate illumination, making it challenging to perform Re-ID by using existing RGB-RGB methods. In practical scenarios, to overcome this limitation, many surveillance cameras can automatically toggle their mode from the visible modality to infrared. This technique enables these cameras to work in dark indoor scenes or at night. Taking advantage of multiple modality cameras, Wu *et al.* [42] introduce a challenging RGB-infrared cross-modality person re-identification task, i.e., *visible-infrared person re-identification* (VI-ReID). Given an infrared (IR) image of a certain person as a probe image, the goal of VI-ReID is to retrieve the corresponding visible (RGB) image of the same person.

Except for the intra-modality variance involved in single-modality ReID, VI-ReID encounters the additional cross-modality discrepancies resulting from the natural difference between the reflectivity

of visible spectrum and the emissivity of thermal spectrum [35]. To tackle these two co-existing problems, we approach the VI-ReID task from an information disentanglement perspective. The foundation of associating visible and infrared images relies on semantic information shared across different modalities and the clues given by this information that establish cross-modality connections. We call this kind of information cross-modality *identity-discriminable information* (IDI), e.g., the shape and outline of a body, and some latent characteristics. Unfortunately, there also exists substantial information only belonging to each modality, e.g., color information for visible images and thermal information for infrared images. We call this kind of information the cross-modality *identity-ambiguous information* (IAI). The retrieval performance will be amortized due to noise caused by the redundant IAI dimensions, as illustrated in Fig. 1. Thus, learned disentanglement representations are supposed to eliminate the IAI and only preserve cross-modality IDI.

Due to the powerful capability of generalization and compaction, variational auto-encoder (VAE) is widely employed to disentanglement representation learning [26]. However, according to [29], conventional VAEs embed multiple classes or data clusters through a standard Gaussian distribution, which is able to model common characteristics of all inputs and restrict the scope of the distribution. For VI-ReID, the standard Gaussian distribution is effective for IAI, but is relatively ineffective for IDI which often incorrectly handles the structural discontinuity between disparate classes in a latent space. Considering that Mixture-of-Gaussians (MoG) model favorably handles the multi-cluster data, embedding IDI by MoG distribution not only explicitly models the intra-class variations but also ensures the inter-class separability.

To this end, we exploit a dual Gaussian-based VAE (DG-VAE), which aims at disentangling the cross-modality feature maps into IDI and IAI codes for the robust VI-ReID task. Specifically, we enforce the IDI codes follow the MoG distribution where each component corresponds to a particular identity. The variance within each component models the intra-identity differences. Different from the IDI codes, the IAI codes are required to follow the normal distribution. Meanwhile, aligning both modalities with the normal distribution is beneficial in regularizing the disentangling process in case reconstruction relies on only the IAI codes. Furthermore, we propose a *triplet swap reconstruction* (TSR) strategy to keep identity-consistency while squeezing IDI and IAI into separate branches, which further promotes disentangling process. To efficiently optimize the proposed DG-VAE, we derive a MoG prior term and a maximum entropy regularizer from maximizing the evidence lower bound (ELBO) based on variational bayesian inference. Considering that our DG-VAE is trained in a supervised manner, we introduce a more powerful adaptive large-margin constraint term to substitute the maximum entropy regularizer, which prevents the MoG distribution to occur “posterior collapse”.

The contributions of this paper are summarized below:

- We present a novel dual Gaussian-based variational disentanglement architecture with the effective triplet swap reconstruction strategy for addressing the VI-ReID task. Such design aims at disentangling identity-discriminable and identity-ambiguous feature subspaces to reduce the modality gap.




Figure 2: Illustrates the proposed probabilistic graphical model for generative process and inference process.

- We theoretically derive the variational inference terms for the proposed MoG prior under supervised setting so that our DG-VAE can be efficiently optimized like standard VAE.
- We experiment with two popular benchmarks where our proposed DG-VAE achieves state-of-the-art performance.

2 RELATED WORK

Visible-infrared person re-identification. Most existing VI-ReID methods could be divided into two groups. The first group [5, 7, 8, 42, 45–47] used only feature-level constraints to reduce the intra- and inter-modality discrepancies, which are similar to RGB-RGB Re-ID methods. For example, Wu *et al.* [42] analyzed three different network structures and used a deep zero padding method for evolving domain-specific structures. Recently, Hao *et al.* proposed a hypersphere manifold embedding [8] and dual alignment embedding [7] to handle intra-modality and inter-modality variations. The second group [38, 39, 41] incorporated generative adversarial network (GAN) to achieve image-level constraints. For instance, Wang *et al.* [38] utilized single-direct alignment strategy, in which IR-RGB retrieval is implemented by matching IR images and fake-IR images generated by corresponding RGB images.

The JSIA [39] is similar to our DG-VAE that both adopt feature disentanglement technique for VI-ReID but it focuses on factorizing modality-invariant and modality-specific information instead of our identity-discriminable and identity-ambiguous information. Moreover, JSIA employed a image stylization architecture to generate cross-modality images by predicting the parameters of AdaIN [12]. *In contrast, our DG-VAE adopts a symmetrical encoder-decoder architecture without AdaIN layer to reconstruct feature maps instead of treating the cross-modality variations as only different styles of images. Empirical results demonstrate our method obtains a better performance than JSIA [39] as shown in Table 1.*

Disentangled representation with VAE. For learning a disentangled representation, one widely used architecture is VAE [19]. Higgins *et al.* [11] introduce an adjustable hyperparameter β to balance the reconstruction and disentanglement quality and propose an unsupervised VAE-based framework named β -VAE. The drawback of β -VAE is that a better disentanglement result is on the expense of worse feature reconstruction. To solve this problem, Kim *et al.* [17] introduce a new penalty that provides a better trade-off.

Afterward, the latent subspace is disentangled based on the specified and unspecified factors [32], which aspect is similar to our work. But it employed only an normal distribution prior instead of a MoG prior. Furthermore, our method reconstructs only the modality-specific features instead of the original images, which reduces the amount of parameters and computational cost. We compare two strategies in Sec. 4.5. The results in Table 3 demonstrate that our method with lightweight decoder still produces the competitive performance compared to image generation methods.

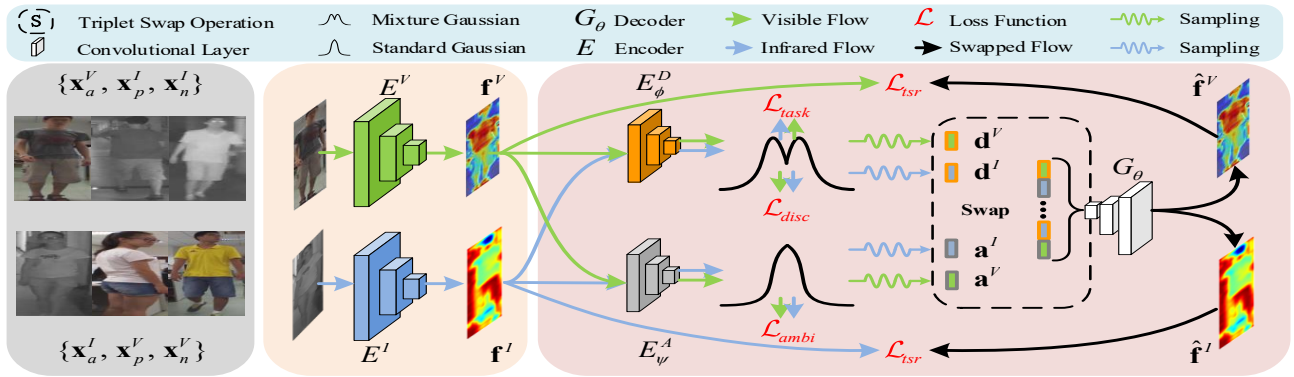


Figure 3: Overview architecture of our proposed DG-VAE.

3 DUAL GAUSSIAN-BASED VAE

Overview. The DG-VAE involves the inference process and generative process (see Fig. 2) whereby we introduce the dual gaussian-based method. First, we adopt the two stream architecture composed of several pre-trained residual blocks as the modality-specific feature extractors, to extract visible-infrared feature maps: $\mathbf{f}_k^V = E^V(\mathbf{X}_k^V)$ and $\mathbf{f}_k^I = E^I(\mathbf{X}_k^I)$, where $\mathbf{f} \in \mathbb{R}^{C \times [H/16] \times [W/16]}$ and C is the number of channels of feature maps. H and W are the height and the width of images \mathbf{X} . Second, to disentangle the IDI and IAI subspaces, we propose dual Gaussian-based priors to constrain the latent codes \mathbf{d} and \mathbf{a} generated by IDI encoder E_ϕ^D and IAI encoder E_ϕ^A , respectively. This results in the IDI code and IAI code conditionally depending on the extracted feature maps, *i.e.*, $\mathbf{d} \sim p_\phi(\mathbf{d}|\mathbf{f})$ and $\mathbf{a} \sim p_\psi(\mathbf{a}|\mathbf{f})$. We apply a latent code classifier to fit the condition probability $p(y|\mathbf{d})$, which enables our DG-VAE to infer the correct identity for VI-ReID task. Third, we design a decoder $p_\theta(\hat{\mathbf{f}}|\mathbf{d}, \mathbf{a})$ with the proposed TSR strategy to promote the disentangling process. To be specific, the reconstructed feature map $\hat{\mathbf{f}}$ is generated by sampling an IAI code \mathbf{a} from $p(\mathbf{a})$ and a corresponding identity y , then, an IDI code \mathbf{d} is sampled from the conditional distribution $\mathbf{d} \sim p(\mathbf{d}|y)$. The decoder $p_\theta(\hat{\mathbf{f}}|\mathbf{d}, \mathbf{a})$ maps the combination of \mathbf{a} and \mathbf{d} to a reconstructed feature map $\hat{\mathbf{f}}$. Note that with $\hat{\mathbf{f}}$, \mathbf{f} , \mathbf{d} and \mathbf{a} without superscripts and subscripts we denote the extracted feature maps and latent codes drawn from both modalities, not distinguishing whether it is from an anchor, positive or negative images. The overall pipeline is shown in Fig. 3.

Based on the above reconstruction process, we employ the idea of variational inference to maximize the ELBO so that the data log-likelihood $\hat{\mathbf{f}}$ is also maximized like the conventional VAE[19]. By using Jensen’s inequality, the log-likelihood in our DG-VAE is:

$$\begin{aligned} \log p(\hat{\mathbf{f}}) &\geq \mathbb{E}_{q_\phi(\mathbf{d}, y|\mathbf{f}), q_\psi(\mathbf{a}|\mathbf{f})} \left[\log \frac{p_\theta(\hat{\mathbf{f}}|\mathbf{d}, \mathbf{a})p(\mathbf{d}, y)p(\mathbf{a})}{q(y|\mathbf{d})q_\phi(\mathbf{d}|\mathbf{f})q_\psi(\mathbf{a}|\mathbf{f})} \right] \\ &= -D_{KL}(q_\phi(\mathbf{d}, y|\mathbf{f})||p(\mathbf{d}, y)) \\ &\quad - D_{KL}(q_\psi(\mathbf{a}|\mathbf{f})||p(\mathbf{a})) \\ &\quad + \mathbb{E}_{q_\phi(\mathbf{d}, y|\mathbf{f}), q_\psi(\mathbf{a}|\mathbf{f})} [\log p_\theta(\hat{\mathbf{f}}|\mathbf{d}, \mathbf{a})]. \end{aligned} \quad (1)$$

Hence, in our method, ELBO consists of three terms in Eq. 1. We call the first term as MoG prior regularizer, which matches $q_\phi(\mathbf{d}, y|\mathbf{f})$ to an identity-specific MoG distribution whose mean and covariance are learned with stochastic gradient variational bayes estimator, and is further introduced in Sec. 3.1. We regard the second

term as standard Gaussian prior regularizer. It pushes $q_\psi(\mathbf{a}|\mathbf{f})$ to align the prior distribution $p(\mathbf{a})$, which is elaborated in Sec. 3.2. The third term is negative reconstruction error, which measures whether the latent code \mathbf{d} and \mathbf{a} are informative enough to recover the original feature maps. In our case, we propose triplet swap reconstruction to achieve this goal, which is described in Sec. 3.3. The whole network is optimized by the multi-objective learning scheme in Sec. 3.4. All further theoretical proof and derivations are elaborated in Appendix A.

3.1 Mixture-of-Gaussians Prior for IDI Encoder

Following the structure of conventional VAEs, we further introduce an identity-discriminable encoder E_ϕ^D upon the above feature extractors (E^V and E^I), to learn IDI representation which enables to identify different persons. Specifically, the IDI codes from visible and infrared images are $\mathbf{d}_k^V = E_\phi^D(\mathbf{f}_k^V)$ and $\mathbf{d}_k^I = E_\phi^D(\mathbf{f}_k^I)$, respectively. These cross-modality IDI codes are organised as multiple clusters on a manifold, where the IDI encoder is supposed to tackle the intra-class variation while pushing inter-class distance.

Unfortunately, conventional VAEs often fail to correctly handle the structural discontinuity between disparate classes in a latent space since they use only a standard Gaussian distribution to embed multiple classes or clusters of data [9, 29, 34]. To solve this problem and further improve the representational capability of IDI codes, \mathbf{d} , we expect them to follow the MoG prior distribution with identity-specific mean μ_y and unit variance. For simplicity, we ignore the correlation among different dimensions of \mathbf{d} , hence the variance is assumed to be diagonal, and the conditional $p(\mathbf{d}|y)$ is thus equal to:

$$p(\mathbf{d}|y) = \mathcal{N}(\mathbf{d}; \mu_y, \mathbf{I}). \quad (2)$$

Recall that in Eq. 1, the KL divergence between $q_\phi(\mathbf{d}, y|\mathbf{f})$ and $p(\mathbf{d}, y)$ is minimized. Since it is difficult to compute the conditional joint probability $q_\phi(\mathbf{d}, y|\mathbf{f})$, a mean-field distribution is used to estimate $q_\phi(\mathbf{d}, y|\mathbf{f})$ under unsupervised conditions in [16]. In contrast, our DG-VAE is trained in a supervised setting and is supposed to maximize the mutual information between a sample and its label. More clarifications are provided in Appendix D. As a consequence, we assume a conditional dependence between \mathbf{d} and the identity y , which results in a variational inference model, $q(\mathbf{d}, y|\mathbf{f})$, illustrated in Fig. 2 and formulated as follows:

$$q(\mathbf{d}, y|\mathbf{f}) = q(y|\mathbf{d})q_\phi(\mathbf{d}|\mathbf{f}). \quad (3)$$

Then, we define the outputs, $\boldsymbol{\mu}_{\phi(\mathbf{f})}$ and $\boldsymbol{\sigma}_{\phi(\mathbf{f})}$, of E_{ϕ}^D as follows:

$$q_{\phi}(\mathbf{d}|\mathbf{f}) = \mathcal{N}(\mathbf{d}; \boldsymbol{\mu}_{\phi(\mathbf{f})}, \boldsymbol{\sigma}_{\phi(\mathbf{f})}^2). \quad (4)$$

Using Eq. 1, 2, 3 and 4, we have:

$$D_{KL}(q_{\phi}(\mathbf{d}, y|\mathbf{f})||p(\mathbf{d}, y)) = \sum_y q(y|\mathbf{d})D_{KL}(q_{\phi}(\mathbf{d}|\mathbf{f})||p(\mathbf{d}|y)) + D_{KL}(q(y|\mathbf{d})||p(y)), \quad (5)$$

where the conditional probability $q(y|\mathbf{d})$ in a supervised setting is the one-hot encoding of the label y . Thus, the first term is to align the latent distribution $q_{\phi}(\mathbf{d}|\mathbf{f})$ and the corresponding class-specific distribution $p(\mathbf{d}|y)$. We call it as MoG prior term. The second term encourages posterior probability $q(y|\mathbf{d})$ to approximate the uniform distribution $p(y)$. In fact, it corresponds to maximize conditional entropy, when the components of the Gaussian are expected to be separable without overlap, which acts as a maximum entropy regularizer. We provide the proof and derivation of the two terms in Appendix B and Appendix C, respectively. If these two terms are exploited in an unsupervised manner, they indeed implement a large-margin clustering algorithm.

Considering that our DG-VAE is a kind of supervised model, the maximum entropy regularizer could be more effectively accomplished by using supervision. Thus, we introduce an adaptive large-margin constrain (ALMC) \mathcal{L}_{lmc} to disperse the components of the MoG distribution with the corresponding learnable margins, inspired by [25]. Finally, by using the reparametrization trick [19] and respectively sampling \mathbf{d} and \mathbf{a} from $q_{\phi}(\mathbf{d}|\mathbf{f})$ and $q_{\psi}(\mathbf{a}|\mathbf{f})$, \mathcal{L}_{gmm} and \mathcal{L}_{lmc} can be calculated by the following closed-form solution:

$$\mathcal{L}_{gmm} = \mathbb{E}_{\mathbf{d} \sim q_{\phi}(\mathbf{d}|\mathbf{f})} \sum_c^{N_y} \mathbf{1}(y=c) \left(\ln(\det(\boldsymbol{\sigma}_{\phi(\mathbf{f})})) + \frac{1}{2} \|\mathbf{d} - \boldsymbol{\mu}_y\|^2 \right), \quad (6)$$

$$\mathcal{L}_{lmc} = -\mathbb{E}_{\mathbf{d} \sim q_{\phi}(\mathbf{d}|\mathbf{f})} \log \frac{\exp^{-D_M(\mathbf{d}, y, \mathbf{f})}}{\sum_{c=1, c \neq y}^{N_y} \exp^{-\|\mathbf{d} - \boldsymbol{\mu}_c\|^2} + \exp^{-D_M(\mathbf{d}, y, \mathbf{f})}}, \quad (7)$$

where N_y is the number of identities in training set. The indicator function $\mathbf{1}()$ equals 1 if y equals c ; and 0 otherwise. The identity-specific means $\boldsymbol{\mu}_y$ are initialized randomly and updated with gradient descent. Furthermore, given some samples following Mixture-of-Gaussians (MoG) distribution, the samples lie on the MoG manifold. The distance between a re-sampled \mathbf{d} with identity y and the corresponding mean $\boldsymbol{\mu}_{\phi(\mathbf{f})}$ should be the squared Mahalanobis distance:

$$D_M(\mathbf{d}, y, \mathbf{f}) = (\mathbf{d} - \boldsymbol{\mu}_{\phi(\mathbf{f})})^T \boldsymbol{\sigma}_{\phi(\mathbf{f})}^{-1} (\mathbf{d} - \boldsymbol{\mu}_{\phi(\mathbf{f})}) - \alpha_y, \quad (8)$$

where the identity-specific margin $\alpha \geq 0$ is a learnable parameter, which can easily achieve the adaptive large-margin MoG constraint, as illuminated in Appendix G.

The combination of \mathcal{L}_{gmm} and \mathcal{L}_{lmc} not only strengthens the network capability of handling multi-class data, but it also constrains the scope of the distribution of each identity by aligning the prior distributions while pushing the inter-class distances. It is robust to outliers or noise so that it has stronger generalization.

The quantitative analysis is elaborated in Sec. 4.5. The objective function of IDI branch is weighted by λ_{gmm} and λ_{lmc} :

$$\mathcal{L}_{disc} = \lambda_{gmm} \mathcal{L}_{gmm} + \lambda_{lmc} \mathcal{L}_{lmc}. \quad (9)$$

3.2 Standard Gaussian Prior for IAI Encoder

To learn the identity-ambiguous representation, we design an IAI encoder E_{ψ}^A to encode the extracted feature maps (\mathbf{f}_{κ}^V and \mathbf{f}_{κ}^I) to IAI codes, which is formulated as: $\mathbf{a}_{\kappa}^V = E_{\psi}^A(\mathbf{f}_{\kappa}^V)$ and $\mathbf{a}_{\kappa}^I = E_{\psi}^A(\mathbf{f}_{\kappa}^I)$.

Unlike the IDI codes that should provide a more reasonable embedding subspace for discriminating different identities, the IAI representation aims at modeling the common characteristics across all the input images. It is scarcely able to identify different persons from these characteristics but they are necessary for feature reconstruction. To this end, we encourage the identity-ambiguous codes to approximate the prior normal distribution $p(\mathbf{a})$ with zero mean and unit variance instead of the MoG prior distribution. We follow [27] to define a Gaussian distribution with mean $\boldsymbol{\mu}_{\psi(\mathbf{f})}$ and diagonal covariance $\boldsymbol{\sigma}_{\psi(\mathbf{f})}$, *i.e.*, $q_{\psi}(\mathbf{a}|\mathbf{f})$, as the output of encoder E_{ψ}^A parametrized by ψ . According to Eq. 1, the ID-ambiguous KL regularization term is formulated as follows,

$$D_{KL}(\mathcal{N}(\boldsymbol{\mu}_{\psi(\mathbf{f})}, \boldsymbol{\sigma}_{\psi(\mathbf{f})}^2) || \mathcal{N}(\mathbf{0}, \mathbf{I})). \quad (10)$$

The KL divergence regularizes the identity-ambiguous codes by limiting the distribution range, such that they do not contain much identity-discriminable information so as to facilitate the disentanglement process [21]. Then, with a reparametrization trick [19], Eq. 10 can be calculated by the following closed solution:

$$\mathcal{L}_{ambi} = \frac{1}{N^V + N^I} \sum_{i=1}^{N^V + N^I} \sum_{l=1}^L \left((\boldsymbol{\mu}_{il})^2 + (\boldsymbol{\sigma}_{il})^2 - \log(\boldsymbol{\sigma}_{il})^2 - 1 \right), \quad (11)$$

where N^V and N^I denote the number of visible and infrared images in a mini-batch, respectively. L is the length (dimension) of the identity-ambiguous code. Eq. 11 allows the latent variational representation to be differentiated and capable of back-propagation.

Furthermore, although we train the encoders separately to extract these information, the decoder G_{θ} may largely rely on the identity-ambiguous codes to synthesize new feature maps, while ignoring the identity-discriminable ones. This situation will distract the feature disentanglement process during training. Thus, using standard Gaussian prior to restrict the scope of feature distribution also enables to effectively mitigate this problem.

3.3 Triplet Swap Reconstruction

Moreover, we propose TSR strategy to perform the whole encoder-decoder process. Different from conventional VAEs, DG-VAE carries out the reconstruction where the decoder uses two types of latent codes swapped across cross-modality triplet as input.

To squeeze the IDI and IAI contained in visible-infrared feature maps into separate branches, we assume that the disentangled representation satisfies the following conditions: 1) An original feature map should be reconstructed from its ID-discriminable and ID-ambiguous code as illustrated in Fig. 4(a); 2) Swapping ID-discriminable codes between cross-modality anchor and positive feature maps, the positive feature maps should be reconstructed from the swapped ID-discriminable code and their ID-ambiguous

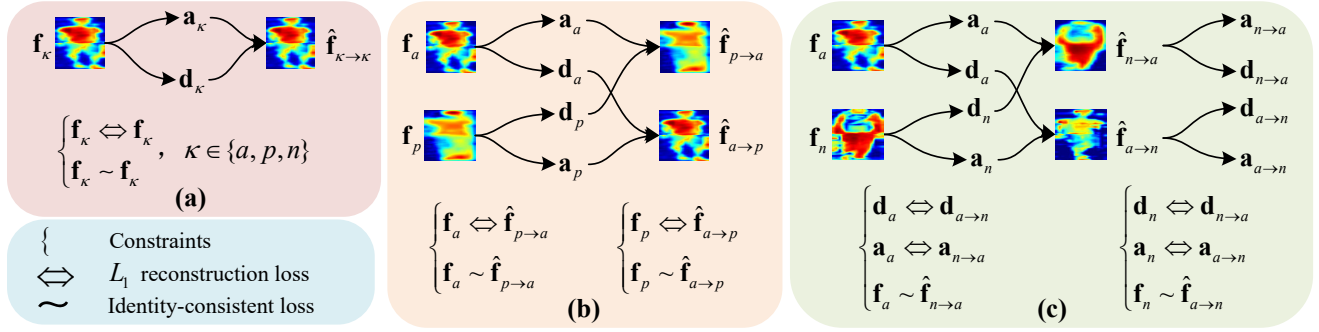


Figure 4: Illustration of the proposed triplet swap reconstruction in detail. For clear representation, we show only the case of the triplet that consists of a visible anchor, infrared positive, and infrared negative. Likewise, the other triplet that contains a infrared anchor, visible positive, and visible negative is swapped and reconstructed.

code, while the identities of reconstructed feature maps should be kept consistently, as shown in Fig. 4(b); 3) Swapping ID-ambiguous codes between a cross-modality pair (whatever positive pair or negative pair), the identities of reconstructed feature maps should correspond to that of the ID-discriminable codes like Fig. 4(c).

To reach these conditions, the decoder G_θ is required to reconstruct an anchor feature map \mathbf{f}_a^V (or \mathbf{f}_a^I) from $\mathbf{d}_a^V \odot \mathbf{a}_a^V$ and $\mathbf{d}_p^I \odot \mathbf{a}_p^V$ (or $\mathbf{d}_a^I \odot \mathbf{a}_a^I$ and $\mathbf{d}_p^V \odot \mathbf{a}_p^I$), where the \odot denotes a concatenation operation as depicted in Fig. 3. Thus, the reconstruction loss is:

$$\mathcal{L}_{rec} = \mathbb{E}_{\kappa \in \{a,p,n\}} \left[\left\| \mathbf{f}_\kappa^V - G_\theta(\mathbf{d}_\kappa^V \odot \mathbf{a}_\kappa^V) \right\|_1 + \left\| \mathbf{f}_\kappa^I - G_\theta(\mathbf{d}_\kappa^I \odot \mathbf{a}_\kappa^I) \right\|_1 \right] + \mathbb{E}_{i,j \in \{a,p\}, i \neq j} \left[\left\| \mathbf{f}_i^V - G_\theta(\mathbf{d}_j^V \odot \mathbf{a}_i^V) \right\|_1 + \left\| \mathbf{f}_i^I - G_\theta(\mathbf{d}_j^I \odot \mathbf{a}_i^I) \right\|_1 \right]. \quad (12)$$

The first term performs a standard self-reconstruction, enforcing the combination of ID-discriminable and ID-ambiguous codes from the same image to contain all information to reconstruct the original features. It is employed as the reconstruction error term in Eq. 1. The second term encourages the encoder E_ϕ^D to extract the shared identity-discriminable features, \mathbf{d}_a^V (or \mathbf{d}_a^I) and \mathbf{d}_p^V (or \mathbf{d}_p^I) from a pair of \mathbf{f}_a^I (or \mathbf{f}_a^V) and \mathbf{f}_p^V (or \mathbf{f}_p^I), focusing on the consistent information between them. Other factors, not shared by \mathbf{f}_a^V and \mathbf{f}_a^I , are encoded into the identity-ambiguous features, \mathbf{d}_a^V (or \mathbf{d}_a^I) and \mathbf{d}_p^I (or \mathbf{d}_p^V). Note that it might be unreasonable to impose an L_1 reconstruction constraint on the reconstructed feature maps generated from cross-modality swapped negative pair, *i.e.*, $G_\theta(\mathbf{d}_a^I \odot \mathbf{a}_n^V)$, $G_\theta(\mathbf{d}_n^I \odot \mathbf{a}_a^V)$, $G_\theta(\mathbf{d}_a^V \odot \mathbf{a}_n^I)$, and $G_\theta(\mathbf{d}_n^V \odot \mathbf{a}_a^I)$. Because both their IDI and IAI are not shared, hence, the reconstructed feature maps are hardly recovered from their combinations.

Furthermore, we introduce the cycle-consistency loss [53] to make the reconstructed features preserve the ID-discriminable and ID-ambiguous codes of their original features. Moreover, given the reconstructed feature $\hat{\mathbf{f}}_\kappa$ ($\hat{\mathbf{f}}_\kappa^V$ or $\hat{\mathbf{f}}_\kappa^I$), the latent code classifier D_C enforces the identity-consistent constraints on all of the codes. The cycle- and identity-consistency loss are defined as follows,

$$\mathcal{L}_{cyc} = \mathbb{E}_{i,j \in \{a,p,n\}} \left[\left\| \mathbf{a}_i - E_\psi^A(G_\theta(\mathbf{d}_j \odot \mathbf{a}_i)) \right\|_1 + \left\| \mathbf{d}_i - E_\phi^D(G_\theta(\mathbf{d}_i \odot \mathbf{a}_j)) \right\|_1 \right], \quad (13)$$

$$\mathcal{L}_{idc} = \mathbb{E}_{\kappa \in \{a,p,n\}} \left[-\log D^C(E_\phi^D(\hat{\mathbf{f}}_\kappa)) \right], \quad (14)$$

where $D^C(E_\phi^D(\hat{\mathbf{f}}_\kappa))$ is the predicted probability of the current feature map $\hat{\mathbf{f}}_\kappa$ belonging to the ground-truth class (identity).

The overall objective for the proposed TSR is given as follows,

$$\mathcal{L}_{tsr} = \lambda_{idc} \mathcal{L}_{idc} + \lambda_{cyc} \mathcal{L}_{cyc} + \lambda_{rec} \mathcal{L}_{rec}, \quad (15)$$

where λ_{idc} , λ_{cyc} and λ_{rec} are the trade-off factors.

3.4 Multi-objective Learning and Optimization

To simultaneously consider the reparameterized latent variate, DG-VAE shares the same feature classifier for both $\boldsymbol{\mu}_\phi(\mathbf{f})$ and \mathbf{d} , which are optimized by the weighted sum of standard cross-entropy loss and cross-modality triplet loss [5], *e.g.*, $\ell(a^V, p^I, n^I) = [D(a^V, p^I) - D(a^V, n^I) + m]_+$ and $\ell(a^I, p^V, n^V) = [D(a^I, p^V) - D(a^I, n^V) + m]_+$, where $D(\cdot, \cdot)$ denotes L_2 distance between two normalized inputs and m is a fixed margin hyperparameter. Given a mini-batch with N^V visible images and N^I infrared images, the retrieval features are formulated as follows,

$$\mathcal{L}_{id} = -\frac{1}{N^V + N^I} \sum_i^{N^V + N^I} \left(\log D^C(\boldsymbol{\mu}_\phi(\mathbf{f}_i)) + \lambda_{aux} \log D^C(\mathbf{d}_i) \right) \quad (16)$$

$$\mathcal{L}_{cmtl} = \frac{1}{N^V} \sum_{(a,p,n)}^{N^V} \left(\ell(\boldsymbol{\mu}_\phi(\mathbf{f}_a^V), \boldsymbol{\mu}_\phi(\mathbf{f}_p^I), \boldsymbol{\mu}_\phi(\mathbf{f}_n^I)) + \lambda_{aux} \ell(\mathbf{d}_a^V, \mathbf{d}_p^I, \mathbf{d}_n^I) \right) + \frac{1}{N^I} \sum_{(a,p,n)}^{N^I} \left(\ell(\boldsymbol{\mu}_\phi(\mathbf{f}_a^I), \boldsymbol{\mu}_\phi(\mathbf{f}_p^V), \boldsymbol{\mu}_\phi(\mathbf{f}_n^V)) + \lambda_{aux} \ell(\mathbf{d}_a^I, \mathbf{d}_p^V, \mathbf{d}_n^V) \right) \quad (17)$$

$$\mathcal{L}_{task} = \lambda_{id} \mathcal{L}_{id} + \lambda_{cmtl} \mathcal{L}_{cmtl}, \quad (18)$$

where $D^C(\cdot)$ is the predicted probability of the input belonging to the ground-truth class. λ_{aux} , λ_{id} and λ_{cmtl} are trade-off factors.

The overall objective is a weighted sum of all loss functions:

$$\mathcal{L}_{all} = \mathcal{L}_{task} + \lambda_{tsr} \mathcal{L}_{tsr} + \lambda_{ambi} \mathcal{L}_{ambi} + \lambda_{disc} \mathcal{L}_{disc} \quad (19)$$

where λ_{tsr} , λ_{ambi} and λ_{disc} are the loss weights to balance each term during the training process detailed in Sec. 4.2. We only use the mean of IDI code, $\boldsymbol{\mu}_\phi(\mathbf{f})$, as retrieval feature at testing time.

4 EXPERIMENTS

4.1 Experimental Setup

Datasets and Evaluation Settings. We evaluated our method on two publicly available datasets: RegDB [33] and SYSU-MM01 [42]. Our experiments followed the RegDB evaluation protocol as described in [46, 47] and the SYSU-MM01 evaluation protocol from [42]. The RegDB dataset consists of 2,060 visible images and 2,060 far-infrared images with 206 identities for training. The testing set contains 206 identities with 2,060 visible images for the query and 2,060 far-infrared images for the gallery. The SYSU dataset contains 22,258 visible images and 11,909 near-infrared images of 395 identities for training. The testing set includes 96 identities with 3,803 near-infrared images for the query and 301 visible images as the gallery. The SYSU dataset is collected by four visible cameras and two near-infrared cameras, used in both indoor and outdoor environments. We adopted the most challenging *single-shot all-search* mode and repeated the above evaluation of 10 trials with a random split of the gallery and probe set to get final results.

Evaluation metrics. We adopt two popular evaluation metrics: rank-k and mean Average Precision (mAP). The rank-k (dubbed R-1 or R-10) indicates the cumulative rate of true matches in the top-k position. The mAP considers person re-identification as a retrieval task, which reflects the comprehensive retrieval performance.

4.2 Implementation Details

DG-VAE is implemented by Pytorch on an NVIDIA Titan Xp GPU.

Mini-batch organization. We follow [7] and resize visible and infrared images to $3 \times 384 \times 128$. Each mini-batch contains 4 different identities, while each identity has 4 pairs of visible and infrared images. Within the mini-batch, we form 32 cross-modality triplets by selecting positive and negative instances for each image, following the rule that the anchor and positive are from the same person in different modalities, while anchor and negative are required to have different identification labels in different modalities.

Network architecture. Since the low-level visual patterns (*e.g.*, texture, contour) of infrared images are similar to general visible images [47], the two modality-specific feature extractors share the same architecture, which consists of three pre-trained residual blocks in [37] with $C = 1024$ channels of output. Note that the parameters of two streams are optimized separately to capture the information of each modality. Moreover, the IDI and IAI encoders respectively contain two pre-trained blocks of ResNet-50[10] followed by two heads for predicting mean and variance. Both heads of IDI encoder is followed by two max-pooling layer, since it favourably captures the discriminative features. The heads of IAI encoder are followed by two avg-pooling layers, since it could provide a comprehensive representation. Meanwhile, the different-modality feature maps share the same IDI and IAI encoder. The main reason is that both modalities include two such types of information and sharing parameters reduces the computational costs. Finally, the decoder consists of a fully connected layer with batch normalization [14], Leaky ReLU [30], Dropout [36] and a series of transposed convolutional layers. It inputs IDI and IAI features whose dimensions are 2048 and 512, respectively. The latent code classifier has only one fully connected layer. The detailed configuration of the architecture is given in Appendix F.

Table 1: Comparison with state-of-the-art VI-ReID methods on the RegDB and the SYSU-MM01 datasets.

	Datasets	RegDB			SYSU-MM01		
	Methods	R-1	R-10	mAP	R-1	R-10	mAP
Feature Extraction	ZERO[42]	17.75	34.21	18.90	14.80	54.12	15.95
	TONE[46]	16.87	34.03	14.92	12.52	50.72	14.42
Metric Learning	HCML[46]	24.44	47.53	20.80	14.32	53.16	16.16
	BCTR[47]	32.67	57.64	30.99	16.12	54.90	19.15
	BDTR[47]	33.47	58.42	31.83	17.01	55.43	19.66
	HSME[8]	41.34	65.21	38.82	18.03	58.31	19.98
	D-HSME[8]	50.85	73.36	47.00	20.68	62.74	23.12
Image Generation	D2RL[41]	43.4	66.10	44.10	28.90	70.60	29.20
	JSIA[39]	48.10	-	48.90	38.10	80.70	36.90
	AlignGAN[38]	56.30	-	53.40	42.40	85.00	40.70
Distribution Alignment	cmGAN[5]	-	-	-	26.97	67.51	27.80
	MAC[45]	36.43	62.36	37.03	33.37	82.49	44.95
	DFE[7]	70.13	86.32	69.14	48.71	88.86	48.59
	XIV[23]	62.21	83.13	60.18	49.92	89.79	50.73
	SSFT[28]	72.30	-	72.90	61.60	89.20	63.20
	DG-VAE	72.97	86.89	71.78	59.49	93.77	58.46

Training strategy. We use the Adam optimizer [18] with $\beta_1 = 0.9$ and $\beta_2 = 0.999$ to train two extractors, two encoders and a decoder, and use stochastic gradient descent with momentum 0.9 for the latent code classifier. Similar to the training scheme in [6], we train DG-VAE in three stages. In the first stage, we train the IDI encoder E_ϕ^D using \mathcal{L}_{task} and \mathcal{L}_{disc} for 100 epochs. A learning rate is set to $2e^{-4}$. In the second stage, we fix the IDI branch, and train the IAI encoder E_ψ^A , the decoder G_θ , and the latent code classifier D^C with the corresponding losses, \mathcal{L}_{tsr} and \mathcal{L}_{ambi} . This process iterates for 50 epochs with learning rate of $2e^{-4}$. Finally, we train the whole network end-to-end with the learning rate of $2e^{-5}$ for 100 epochs. We augment the datasets with horizontal flipping and random erasing [52].

Hyperparameter. Following the parameters settings in [48] and [5], we set $\lambda_{aux} = 0.1$ and $m = 0.1$. In the first stage, we empirically find that training with a large value of λ_{ambi} is unstable. We thus set λ_{ambi} to 0.001 in the second stage, and increase it to 0.01 in the third stage to regularize the disentanglement. We fix λ_{id} , λ_{cmtl} , λ_{idc} and λ_{cyc} to 1, 1, 0.5 and 0.5, respectively. For other parameters, we fix the split of RegDB dataset and use corresponding images as training/validation sets. We use a grid search on the validation split to set the parameters, resulting in $\lambda_{gmm} = 1$, $\lambda_{lmc} = 0.1$, $\lambda_{disc} = 0.2$ and $\lambda_{rec} = 0.5$. We fix all parameters on both datasets.

4.3 Comparison with State-of-the-art Methods

Our proposed DG-VAE outperforms the four types of state-of-the-art VI-ReID methods (see Table 1). On the one hand, DG-VAE could be treated as a distribution alignment method since we enforce both modalities to follow the same priors. On the other hand, our proposed method indeed contains the encoder-decoder architecture but does not generate original images. Hence, we analyze two of the most related types of methods as follows.

Image Generation. Both D2RL[41] and JSIA[39] employ image stylization networks to generate cross-modality images by predicting the parameters of the AdaIN layers [12], which are shown to

Table 2: Evaluation of our baseline and its variants with different loss functions on the SYSU-MM01 and RegDB dataset.

	Loss Functions								SYSU-MM01		RegDB	
	\mathcal{L}_{id}	\mathcal{L}_{cmtl}	\mathcal{L}_{rec}	\mathcal{L}_{idc}	\mathcal{L}_{cyc}	\mathcal{L}_{ambi}	\mathcal{L}_{gmm}	\mathcal{L}_{lmc}	Rank-1	mAP	Rank-1	mAP
Baseline	✓	✓							44.45	44.89	58.74	59.80
Mixture-of-Gaussians Prior	✓	✓					✓		42.12	42.37	58.12	59.24
	✓	✓					✓	✓	45.39	45.41	60.28	60.92
Two-branch AE with TSR	✓	✓	✓						48.91	48.22	62.67	63.16
	✓	✓	✓	✓					49.02	48.69	62.81	63.47
	✓	✓	✓	✓	✓				49.01	48.93	62.94	63.42
Two-branch VAE with TSR (one standard Gaussian for IAI branch)	✓	✓	✓				✓		54.43	53.83	68.20	67.92
	✓	✓	✓	✓			✓		55.97	54.21	68.74	68.83
	✓	✓	✓	✓	✓		✓		56.65	55.82	69.13	69.01
Two-branch VAE with TSR (two standard Gaussians for two branches)	✓	✓	✓	✓	✓	✓			<u>57.08</u>	<u>56.46</u>	<u>69.84</u>	<u>69.75</u>
DG-VAE with TSR (one standard Gaussian and one MoG)	✓	✓	✓	✓	✓	✓	✓		56.13	55.27	68.49	68.20
	✓	✓	✓	✓	✓	✓	✓	✓	59.49	58.46	72.97	71.78

mainly captures the style information of the image. However, the visible-infrared discrepancy is more complex than the variation of style characteristics, e.g., unaligned pose variation especially in the more challenging SYSU-MM01 dataset. In contrast, our DG-VAE and AlignGAN[38] utilize the encoder-decoder network without AdaIN layer to model such an intractable situation and reach a better performance. However, without a prior assumption, AlignGAN directly maps the learned latent codes to generated images, which leads model to learn only a one-to-one mapping. The AlignGAN struggles to generate cross-modality images when it encounters the unseen inputs in test set, thereby degrading the generalization. Unlike the above methods, DG-VAE samples latent codes from the proposed dual Gaussian priors to reconstruct cross-modality feature maps for disentangling IDI and IAI, which results in the significant improvements compared with AlignGAN [38], rank-1 accuracy by 16.67% and mAP by 18.38% on the RegDB dataset.

Distribution Alignment. The DFE [7] achieves the most competitive results, which indicates that bridging the cross-modality gap by decreasing the distribution divergence is effective. However, they estimated a bias distribution drawn from only the observed data. Benefiting from the reparametrization trick, DG-VAE takes additional latent data from re-sampling operation into account, which allows model to explore unobserved data and approximates the estimated distribution with true distribution. Hence, our DG-VAE outperforms DFE [7] in terms of rank-1 accuracy by 10.78% and mAP by 9.87% on the SYSU-MM01 dataset.

4.4 Ablation Study

The first row in Table 2 is the baseline model composed of E^V , E^I and E_ϕ^D , which is optimized by \mathcal{L}_{task} with $\lambda_{aux} = 0$.

Impact of ALMC. To verify the effectiveness of ALMC, we reset λ_{aux} to 0.1. The first three rows indicate that adding only the MoG prior term \mathcal{L}_{gmm} to the baseline model will degrade performance. It is because at the beginning the distribution of each cluster is not stable. Meanwhile, each initial Gaussian component is stochastic, which has a high probability of becoming degenerated. Therefore, when we impose the ALMC \mathcal{L}_{lmc} to collaboratively optimize the

model, the performance on both datasets increases with a considerable gain.

Impact of TSR. To demonstrate the effectiveness of TSR, based on the baseline model, we further build a two-branch AE by adding the decoder G_θ and the IAI encoder E_ψ^A and directly use the separate means of IDI and IAI codes for reconstruction without reparameterization operation. The fourth row shows that the TSR strategy improves the performance of baseline even with the general AE architecture. However, applying the identity and cycle consistent constraints (\mathcal{L}_{idc} and \mathcal{L}_{cyc}) is not efficient in such AE. This is because without the augmented samples from the re-sampling operation, the identity-consistent condition is relatively easy to achieve by the reconstruction loss L_{rec} .

Impact of dual Gaussian priors. To comprehensively explore the advantages of the proposed dual Gaussian priors, we conduct three series of experiments to verify the performances of two-branch VAEs with one standard Gaussian, two standard Gaussian and the combining of one MoG and one standard Gaussian priors. First, we employ our full model while only executing re-sampling operation on IAI branch. The seventh row shows that such architecture with TSR still obtains improvements on two datasets. The eighth and ninth rows indicate that applying the \mathcal{L}_{idc} and \mathcal{L}_{cyc} further boosts the VI-ReID performance with the VAE architecture. Second, we adopt full model architecture but calculate \mathcal{L}_{ambi} on both branches with re-sampling operation. The eleventh row shows using two standard Gaussian priors obtains mere improvement. Although restricting the IDI following standard Gaussian distribution is beneficial for learning compacted representations, only a standard Gaussian distribution is hard to model the complex multi-cluster structures. Finally, the last two rows demonstrate that our MoG prior effectively alleviates this drawback. It implies that the MoG prior and the triplet swap disentangling are complementary, and combining all the loss terms produces the best results.

4.5 Discussion

Why are MoG prior term and ALMC efficient? One advantage of the DG-VAE is to represent feature embedding as a distribution instead of a fixed feature vector. When our model encounters the

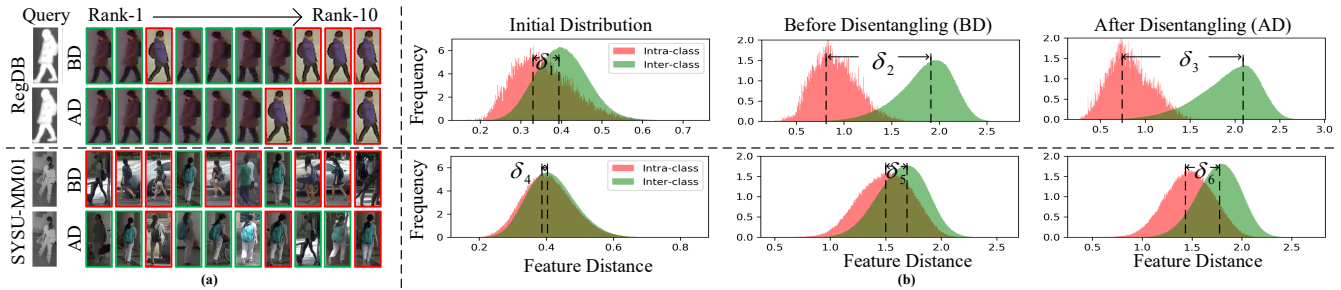


Figure 5: (a) is the illustration of the qualitative examples. The green borders donate the correct matches corresponding the given query and the red is opposite. (b) depicts the distributions of the Euclidean distance between cross-modality (RGB-IR) features. The intra-class and inter-class distances are indicated by red and green color, respectively.

Table 3: Comparison of different reconstruction strategies.

Datasets	RegDB		Decoder & Discriminator	
Strategies	R-1	mAP	Params(M)	FLOPs(G)
Feature Maps	72.97	71.78	16.26	1.23
<i>with discriminator</i>	73.41	72.09	18.52	1.50
Original Images	72.62	72.15	31.22	6.03
<i>with discriminator</i>	73.74	72.85	47.10	6.56

outliers caused by noise label and occlusion as shown in Appendix E, it will assign the larger variance to these samples instead of sacrificing inter-class separability, thereby promoting robust feature learning. In addition, the last three rows of Table 2 suggest that the MoG prior is more beneficial for feature disentanglement than feature learning. As for feature disentanglement, class-specific prior leads the model to explicitly handle the variation of each identity, which is equivalent to extracting IDI across different modalities. It allows this kind of information to flow into the IDI branch, thereby promoting the disentanglement process. On the other hand, a more reasonable distribution of latent variates enables more realistic yields from the generator. In short, with the ALMC, the MoG prior makes extracting the disentangled features more efficient.

Why is triplet swap disentanglement efficient for VI-ReID?

The retrieval features learned by most existing VI-ReID methods are often highly entangled representations. Intuitively, cross-modality tasks suffer from more data discrepancies than single-modality tasks. Thus, excluding IAI plays a crucial role in boosting performance. The qualitative examples are shown in Fig. 5(a). Another reason for its efficiency becomes clear if we count the frequency of inter-class and intra-class distance as illustrated in Fig. 5(b). Comparing the first and second columns, the means of inter- and intra-distance are pushed away by using the baseline method with the MoG prior, where $\delta_1 < \delta_2$ and $\delta_4 < \delta_5$. Note that due to the difficulty of the SYSU-MM01 dataset, where $\delta_4 \approx 0$, it is hard to significantly push two peaks away like for the RegDB dataset, where $\delta_2 > \delta_5$. Thanks to disentangling ID-discriminable factors, our full model performs considerably better on both datasets, $\delta_3 \approx 1.4$ and $\delta_6 \approx 0.4$, which means 13.05% and 10.86% improvement of mAP on the SYSU-MM01 and RegDB datasets, respectively. As a consequence, based on factorized IDI, DG-VAE indeed leads the

intra-class embeddings to be compacted and disperses the inter-class clusters with a large margin as illustrated in Appendix G, thereby achieving better performance.

Why does DG-VAE reconstruct the feature map instead of the original image? We experiment with different settings as shown in Table 3. Our full DG-VAE model is regarded as baseline setting in the first row. As shown in the second row, for a fair comparison, we treat the reconstructed feature maps as a multi-channel image and add an advanced PatchGAN discriminator [15] to distinguish whether the generated multi-channel image is fake. Moreover, we add four extra deconvolutional layers to generate fake images with the same height and width as the original image, but don't calculate binary cross-entropy loss function of the discriminator. The difference between the third and fourth rows is whether the binary cross-entropy loss function of the discriminator is calculated or not. The results show that our feature maps reconstruction strategy gets a competitive performance compared to generating the original image. Meanwhile, our model requires only half of the parameters and one-fifth FLOPs. We find increased performance can be attributed to the extra discriminator with binary cross-entropy loss function. Nevertheless, we do not have the motivation to distinguish reconstructed feature maps. On the other hand, DG-VAE starts from the dynamically learned feature maps instead of the fixed input images, which is different from the assumption of Bayes evidential reasoning.

5 CONCLUSION

In this paper, we design a novel variational disentanglement architecture with dual Gaussian constraints and triplet swap reconstruction for robust cross-modality VI-ReID. By excluding ID-ambiguous information, we show that the learned representations achieve significant improvement compared to the baseline and reaches a competitive performance with state-of-the-art methods. Furthermore, we analyze the reason for the observed increase in performance, and discuss the principle of our derived MoG prior regularizer. The proposed method shows considerable potential of the disentangled representation for multi-modality tasks.

ACKNOWLEDGMENTS

This work was supported mainly by the LIACS Media Lab at Leiden University and in part by the China Scholarship Council. We are also grateful to the support of NVIDIA with the grant of one Titan Xp GPU card.

REFERENCES

- [1] Alexander A. Alemi, Ian Fischer, Joshua V. Dillon, and Kevin Murphy. 2017. Deep Variational Information Bottleneck. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24–26, 2017, Conference Track Proceedings*.
- [2] Binghui Chen, Weihong Deng, and Jiani Hu. 2019. Mixed high-order attention network for person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*. 371–381.
- [3] Dapeng Chen, Dan Xu, Hongsheng Li, Nicu Sebe, and Xiaogang Wang. 2018. Group consistent similarity learning via deep crf for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 8649–8658.
- [4] Tianlong Chen, Shaojin Ding, Jingyi Xie, Ye Yuan, Wuyang Chen, Yang Yang, Zhou Ren, and Zhangyang Wang. 2019. ABD-Net: Attentive but Diverse Person Re-Identification. *arXiv preprint arXiv:1908.01114* (2019).
- [5] Pingyang Dai, Rongrong Ji, Haibin Wang, Qiong Wu, and Yuyu Huang. 2018. Cross-Modality Person Re-Identification with Generative Adversarial Training. In *IJCAI*. 677–683.
- [6] Yixiao Ge, Zhuowan Li, Haiyu Zhao, Guojun Yin, Shuai Yi, Xiaogang Wang, et al. 2018. Fd-gan: Pose-guided feature distilling gan for robust person re-identification. In *Advances in neural information processing systems*. 1222–1233.
- [7] Yi Hao, Nannan Wang, Xinbo Gao, Jie Li, and Xiaoyu Wang. 2019. Dual-alignment Feature Embedding for Cross-modality Person Re-identification. In *Proceedings of the 27th ACM International Conference on Multimedia*. 57–65.
- [8] Yi Hao, Nannan Wang, Jie Li, and Xinbo Gao. 2019. HSME: hypersphere manifold embedding for visible thermal person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 8385–8392.
- [9] Junxian He, Daniel Spokoyny, Graham Neubig, and Taylor Berg-Kirkpatrick. 2019. Lagging inference networks and posterior collapse in variational autoencoders. *arXiv preprint arXiv:1901.05534* (2019).
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [11] Irina Higgins, Loïc Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. 2017. beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24–26, 2017, Conference Track Proceedings*.
- [12] Xun Huang and Serge J. Belongie. 2017. Arbitrary Style Transfer in Real-Time with Adaptive Instance Normalization. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22–29, 2017*. IEEE Computer Society, 1510–1519.
- [13] Yan Huang, Qiang Wu, JingSong Xu, and Yi Zhong. 2019. Sbsgan: Suppression of inter-domain background shift for person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*. 9527–9536.
- [14] Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167* (2015).
- [15] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1125–1134.
- [16] Zhuxi Jiang, Yin Zheng, Huachun Tan, Bangsheng Tang, and Hanning Zhou. 2017. Variational Deep Embedding: An Unsupervised and Generative Approach to Clustering. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19–25, 2017*, Carles Sierra (Ed.), ijcai.org, 1965–1972.
- [17] Hyunjik Kim and Andriy Mnih. 2018. Disentangling by Factorising. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10–15, 2018 (Proceedings of Machine Learning Research)*, Jennifer G. Dy and Andreas Krause (Eds.), Vol. 80. PMLR, 2654–2663.
- [18] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [19] Diederik P. Kingma and Max Welling. 2014. Auto-Encoding Variational Bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14–16, 2014, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.).
- [20] Hussam Lawen, Avi Ben-Cohen, Matan Protter, Itamar Friedman, and Lihi Zelnik-Manor. 2019. Attention Network Robustification for Person ReID. *arXiv preprint arXiv:1910.07038* (2019).
- [21] Hsin-Ying Lee, Hung-Yu Tseng, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. 2018. Diverse image-to-image translation via disentangled representations. In *Proceedings of the European conference on computer vision (ECCV)*. 35–51.
- [22] Qingming Leng, Mang Ye, and Qi Tian. 2019. A survey of open-world person re-identification. *IEEE Transactions on Circuits and Systems for Video Technology* (2019).
- [23] Diangang Li, Xing Wei, Xiaopeng Hong, and Yihong Gong. 2020. Infrared-Visible Cross-Modal Person Re-Identification with an X Modality. In *AAAI*. 4610–4617.
- [24] Yu-Jhe Li, Ci-Siang Lin, Yan-Bo Lin, and Yu-Chiang Frank Wang. 2019. Cross-dataset person re-identification via unsupervised pose disentanglement and adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*. 7919–7929.
- [25] Hao Liu, Xiangyu Zhu, Zhen Lei, and Stan Z Li. 2019. Adaptiveface: Adaptive margin and sampling for face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 11947–11956.
- [26] Yu Liu, Fangyin Wei, Jing Shao, Lu Sheng, Junjie Yan, and Xiaogang Wang. 2018. Exploring disentangled feature representation beyond face identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2080–2089.
- [27] Boyu Lu, Jun-Cheng Chen, and Rama Chellappa. 2019. Unsupervised domain-specific deblurring via disentangled representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 10225–10234.
- [28] Yan Lu, Yue Wu, Bin Liu, Tianzhu Zhang, Baopu Li, Qi Chu, and Nenghai Yu. 2020. Cross-modality Person re-identification with Shared-Specific Feature Transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 13379–13389.
- [29] James Lucas, George Tucker, Roger B Grosse, and Mohammad Norouzi. 2019. Don't Blame the ELBO! A Linear VAE Perspective on Posterior Collapse. In *Advances in Neural Information Processing Systems*. 9403–9413.
- [30] Andrew L Maas, Awni Y Hannun, and Andrew Y Ng. 2013. Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml*, Vol. 30. 3.
- [31] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, Nov (2008), 2579–2605.
- [32] Michael F Mathieu, Junbo Jake Zhao, Junbo Zhao, Aditya Ramesh, Pablo Sprechmann, and Yann LeCun. 2016. Disentangling factors of variation in deep representation using adversarial training. In *Advances in neural information processing systems*. 5040–5048.
- [33] Dat Nguyen, Hyung Hong, Ki Kim, and Kang Park. 2017. Person recognition system based on a combination of body images from visible light and thermal cameras. *Sensors* 17, 3 (2017), 605.
- [34] Ali Razavi, Aaron van den Oord, Ben Poole, and Oriol Vinyals. 2019. Preventing Posterior Collapse with delta-VAEs. In *International Conference on Learning Representations*.
- [35] M Saquib Sarfraz and Rainer Stiefel. 2017. Deep perceptual mapping for cross-modal face recognition. *International Journal of Computer Vision* 122, 3 (2017), 426–438.
- [36] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15, 1 (2014), 1929–1958.
- [37] Guanshuo Wang, Yufeng Yuan, Xiong Chen, Jiwei Li, and Xi Zhou. 2018. Learning discriminative features with multiple granularities for person re-identification. In *Proceedings of the 26th ACM international conference on Multimedia*. 274–282.
- [38] Guan'an Wang, Tianzhu Zhang, Jian Cheng, Si Liu, Yang Yang, and Zengguang Hou. 2019. RGB-Infrared Cross-Modality Person Re-Identification via Joint Pixel and Feature Alignment. In *Proceedings of the IEEE International Conference on Computer Vision*. 3623–3632.
- [39] Guan-An Wang, Tianzhu Zhang Yang, Jian Cheng, Jianlong Chang, Xu Liang, Zengguang Hou, et al. 2020. Cross-Modality Paired-Images Generation for RGB-Infrared Person Re-Identification. *arXiv preprint arXiv:2002.04114* (2020).
- [40] Zheng Wang, Zhixiang Wang, Yang Wu, Jingdong Wang, and Shin'ichi Satoh. 2019. Beyond Intra-modality Discrepancy: A Comprehensive Survey of Heterogeneous Person Re-identification. *arXiv preprint arXiv:1905.10048* (2019).
- [41] Zhixiang Wang, Zheng Wang, Yinqiang Zheng, Yung-Yu Chuang, and Shin'ichi Satoh. 2019. Learning to Reduce Dual-level Discrepancy for Infrared-Visible Person Re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Vol. 2. 4.
- [42] Ancong Wu, Wei-Shi Zheng, Hong-Xing Yu, Shaogang Gong, and Jianhuang Lai. 2017. RGB-Infrared Cross-Modality Person Re-Identification. In *The IEEE International Conference on Computer Vision (ICCV)*.
- [43] Jinlin Wu, Yang Yang, Hao Liu, Shengcai Liao, Zhen Lei, and Stan Z Li. 2019. Unsupervised Graph Association for Person Re-Identification. In *Proceedings of the IEEE International Conference on Computer Vision*. 8321–8330.
- [44] Bryan Ning Xia, Yuan Gong, Yizhe Zhang, and Christian Poellabauer. 2019. Second-Order Non-Local Attention Networks for Person Re-Identification. In *Proceedings of the IEEE International Conference on Computer Vision*. 3760–3769.
- [45] Mang Ye, Xiangyuan Lan, and Qingming Leng. 2019. Modality-aware Collaborative Learning for Visible Thermal Person Re-Identification. In *Proceedings of the 27th ACM International Conference on Multimedia*. 347–355.
- [46] Mang Ye, Xiangyuan Lan, Jiawei Li, and Pong C Yuen. 2018. Hierarchical discriminative learning for visible thermal person re-identification. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [47] Mang Ye, Zheng Wang, Xiangyuan Lan, and Pong C Yuen. 2018. Visible Thermal Person Re-Identification via Dual-Constrained Top-Ranking. In *IJCAI*. 1092–1099.
- [48] Tianyuan Yu, Da Li, Yongxin Yang, Timothy M Hospedales, and Tao Xiang. 2019. Robust Person Re-Identification by Modelling Feature Uncertainty. In *Proceedings*

- of the *IEEE International Conference on Computer Vision*. 552–561.
- [49] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. 2015. Scalable person re-identification: A benchmark. In *Proceedings of the IEEE international conference on computer vision*. 1116–1124.
- [50] Liang Zheng, Yi Yang, and Alexander G Hauptmann. 2016. Person re-identification: Past, present and future. *arXiv preprint arXiv:1610.02984* (2016).
- [51] Zhedong Zheng, Xiaodong Yang, Zhiding Yu, Liang Zheng, Yi Yang, and Jan Kautz. 2019. Joint discriminative and generative learning for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2138–2147.
- [52] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. 2017. Random erasing data augmentation. *arXiv preprint arXiv:1708.04896* (2017).
- [53] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*. 2223–2232.