

# Rethinking Multi-view Representation Learning via Distilled Disentangling

Guangzhou Ke<sup>1,\*</sup>, Bo Wang<sup>2,†</sup>, Xiaoli Wang<sup>3</sup>, Shengfeng He<sup>4,†</sup>

<sup>1</sup>Beijing Jiaotong University,

<sup>2</sup>State Key Laboratory of Multimodal Artificial Intelligence Systems,  
Institute of Automation, Chinese Academy of Sciences,

<sup>3</sup>Nanjing University of Science and Technology, <sup>4</sup>Singapore Management University

guangzhouk@gmail.com, wangbo@ia.ac.cn, xiaoliwang@njust.edu.cn, shengfenghe@smu.edu.sg

## Abstract

Multi-view representation learning aims to derive robust representations that are both view-consistent and view-specific from diverse data sources. This paper presents an in-depth analysis of existing approaches in this domain, highlighting a commonly overlooked aspect: the redundancy between view-consistent and view-specific representations. To this end, we propose an innovative framework for multi-view representation learning, which incorporates a technique we term ‘distilled disentangling’. Our method introduces the concept of masked cross-view prediction, enabling the extraction of compact, high-quality view-consistent representations from various sources without incurring extra computational overhead. Additionally, we develop a distilled disentangling module that efficiently filters out consistency-related information from multi-view representations, resulting in purer view-specific representations. This approach significantly reduces redundancy between view-consistent and view-specific representations, enhancing the overall efficiency of the learning process. Our empirical evaluations reveal that higher mask ratios substantially improve the quality of view-consistent representations. Moreover, we find that reducing the dimensionality of view-consistent representations relative to that of view-specific representations further refines the quality of the combined representations. Our code is accessible at: <https://github.com/Guangzhou-Ke/MRDD>.

## 1. Introduction

Multi-view representation learning (MvRL) [46] forms the cornerstone of various multi-view applications, such as video understanding [7, 17], 3D rendering [62], and cross-

\*This work was done while interning at the Institute of Automation, Chinese Academy of Sciences.

†Corresponding authors.

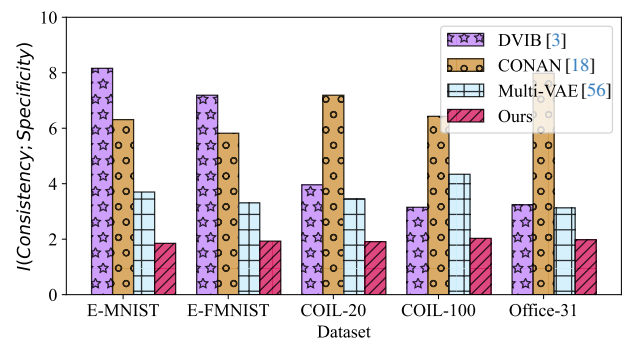


Figure 1. Existing multi-view representation learning methods show high inter-view correlations. We estimate the mutual information of multi-view consistency and specificity of three baseline MvRL models DVIB [3], CONAN [18], Multi-VAE [56], and our method using MINE [4] on the same settings across five datasets.

modal retrieval [35]. In the MvRL context, ‘‘views’’ commonly refer to distinct angles from which objects are captured by cameras or data descriptors, like the histogram of oriented gradients (HOG) [9] and the scale-invariant feature transform (SIFT) [32]. The success of multi-view applications relies on effectively leveraging shared information (*consistency*) among views and distinctive information (*specificity*) within each view. However, learning high-quality view-consistent and view-specific representations from multiple sources poses an open challenge.

Optimal view-consistent and view-specific representations should exhibit both robust representational abilities and minimal redundancy. Reducing redundancy between these two aspects not only improves the quality of combined representations but also decreases the computational burden for subsequent tasks. However, prevailing state-of-the-art methods [3, 11, 22, 56, 57] often neglect the critical aspect of minimizing redundancy between consistency and specificity. As demonstrated in Fig. 1, there is a notable correlation among representations derived from exist-

ing end-to-end approaches<sup>1</sup>. This leads us to pose an important question: *What factors contribute to the dependency within multi-view representations?* In this paper, we address this question from a disentanglement perspective:

**(i) In the unsupervised setting, the joint learning (or end-to-end) paradigm presents significant challenges.**

Within the scope of disentangling representations, most end-to-end MvRL methods aim to extract view-consistent representations by maximizing their mutual information lower bound across views, while simultaneously minimizing the mutual information upper bound between views to derive view-specific representations. This approach essentially forms a Min-Max game, posing a risk for models to settle on suboptimal solutions in the absence of supplementary information. Recent advancements have sought to refine the model’s proficiency in learning high-quality representations by incorporating auxiliary constraints into the joint loss function, such as adversarial [61] and contrastive constraints [18, 21, 42, 49]. Although these approaches mitigate some limitations of joint representation learning, they often overlook a critical factor: the model’s initial inability to differentiate between view-consistent and view-specific information, leading to the accumulation of redundancy.

We propose that if a model could effectively strip away information linked to pre-existing knowledge from multi-view representations, the remaining data would be devoid of such knowledge. We term this process ‘distilled disentangling.’ In our methodology, view-consistent representations are treated as prior knowledge, under the rationale that such consistency represents information common to all views and remains unchanged irrespective of the view. By identifying and excluding the view-consistent information, the model can more accurately isolate view-specific representations for each view.

**(ii) Disparity in Information Density Between Consistency and Specificity.** Extracting view-specific information typically involves processing individual views, whereas view-consistent information necessitates the integration of data from all views. Crafting a unified representation from multiple sources is complex and can lead to a significant escalation in computational resources as the number of views increases. To counter this challenge, some approaches employ multiple lightweight view-specific autoencoders to generate multi-view latent representations [56, 57], followed by an exploration of consistency using these synthesized representations. However, each latent representation is heavily laden with view-specific information, posing a challenge in mitigating this interference.

Our solution introduces the concept of *masked cross-view prediction (MCP)*, which facilitates the learning of multi-view consistency using a single consistent encoder, without the need for additional computational resources.

This is achieved by selectively masking parts of the content and prompting the encoder to predict the masked content by synthesizing visible portions from multi-view data. The advantages of MCP are manifold: 1) it efficiently processes all view data without escalating computational demands; 2) the randomness of masking aids in minimizing the impact of view-specific information; and 3) it strengthens the resilience of view-consistent representations.

Building on this analysis, we adeptly tackle the identified challenges. Leveraging the Masked Cross-View Prediction (MCP) strategy, we employ a single consistent encoder to process all views, yielding high-quality view-consistent representations. Initially, all unmasked blocks are fed into the consistent encoder simultaneously to derive these representations. This step is followed by utilizing multiple decoders to predict the masked content in their respective views, using the same view-consistent representation. The primary benefit of this method is its efficiency in extracting concise consistent representations while concurrently minimizing computational demands. Our experiments demonstrate that, even with a high mask ratio, for instance, 80%, the MCP strategy enables the model to learn superior representations compared to the no-MCP baseline.

To address the issue (i), we freeze the consistent encoder post-training, maintaining the invariance of view-consistent representations during the distillation process for view-specific representations. Additionally, we design a disentangling module that minimizes the upper bound of mutual information between the consistent and view-specific representations, thereby extracting refined view-specific representations. To prevent trivial solutions in the disentangling module, we concatenate the consistent and view-specific representations, employing view-specific decoders to reconstruct the original data. These integrated approaches culminate in our novel multi-view representation learning method, which we term Multi-View Representation learning via Distilled Disentangling (MRDD). The key contributions of this paper are outlined as follows:

- We illuminate fundamental challenges in multi-view representation learning through a disentanglement lens, revealing how these limitations impede the effectiveness of existing models.
- We introduce a multi-view representation learning framework centered on distilled disentangling, which offers a fresh perspective on crafting low-redundancy view-consistent and view-specific representations. Our extensive experimental analysis confirms the superiority of our approach over current state-of-the-art methods.
- Our experimental findings highlight two key insights: i) a high masked ratio (e.g., 80%) significantly enhances the quality of consistent representations; ii) reducing the dimensionality of consistent representations relative to specific representations markedly boosts the quality of their

<sup>1</sup>More details are included in the supplementary materials.

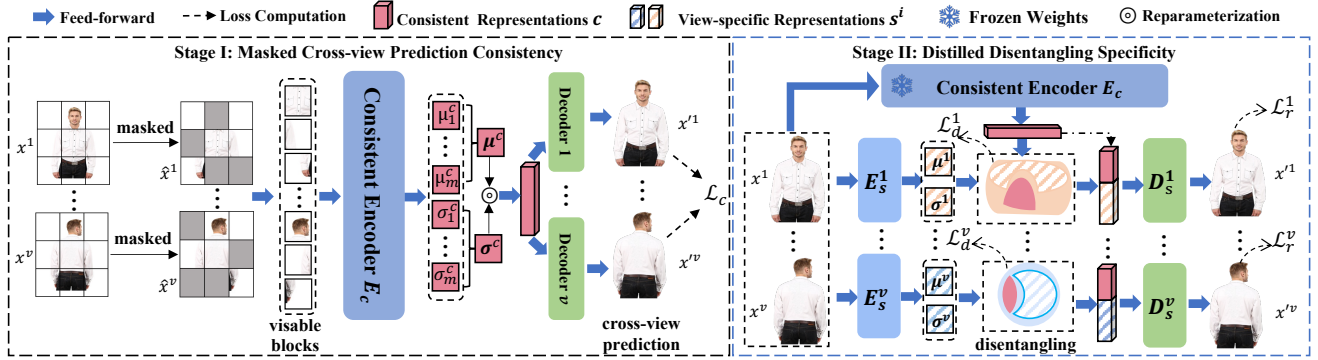


Figure 2. Illustration of the workflow of the proposed framework. The objective of stage I is to exploit the masked cross-view prediction strategy to uncover view-consistent representations. Initially, a consistent encoder is employed to learn consistent representations from all masked data. Ultimately, several decoders are utilized to predict the removed content in the corresponding views. The objective of stage II is to obtain high-quality view-specific representations by filtering out consistency-related information in specific representations. We assume the standard Gaussian distribution as the prior for all representations.

combined representations. We believe that these discoveries will inspire further research in the field of MvRL.

## 2. Related Work

**Multi-view Representation Learning.** The goal of MvRL is to extract both shared and view-specific information from multiple data sources, integrating them into a cohesive representation that is advantageous for predictive tasks [6, 19, 27, 39]. Existing approaches in this field generally fall into three categories: statistic-based [9, 29, 29, 35, 50, 50], deep learning-based [1, 18, 42, 46, 55, 60], and hybrid methods [16, 40, 47]. More discussions are given in the supplementary materials.

Our approach is categorized under deep learning-based methods. We distinguish our work by utilizing deep learning’s capacity to handle large datasets effectively. Moreover, we address the interpretability challenges in representations by incorporating disentanglement techniques

**Disentangled Representation Learning.** Disentangled representation learning (DRL) aims to unravel independent control factors within latent representations [5, 36, 37, 53, 54]. As DRL encompasses a vast area, detailed insights can be found in a comprehensive survey by Wang et al. [48]. Here, we concentrate on DRL specifically within the realm of multi-view learning. Multi-view DRL can be broadly classified into two streams: model-based disentangled methods [18, 41, 42, 57] and information-theoretic disentangled methods [3, 11, 22, 56].

Model-based disentangled methods aim to compel the model to differentiate between shared and view-specific information, often employing constraints like adversarial [61] and contrastive constraints [18, 20]. For instance, Zhang et al. [60] developed a framework with nested autoencoders: the inner autoencoders focus on view-specific information, while the outer ones consolidate this information to extract globally shared knowledge. However, these methods

are often constrained by their reliance on specific model architectures. Information-theoretic disentangled methods, on the other hand, strive for generalization. A representative work [3] utilizes the information bottleneck principle, predicting each view’s data using information from other views.

Our approach employs a distilled disentangling strategy, initially extracting consistent representations across views and subsequently decoupling each view to derive view-specific representations.

**Masked Modeling.** The masked modeling technique is a self-supervised learning approach that involves predicting masked components based on the context. In natural language processing, BERT [23] employs masked sequences to enhance contextual understanding, a technique now widespread in training large language models [2, 26, 31]. In the visual domain, denoising autoencoders [45] use a similar principle to learn robust representations by reconstructing corrupted parts of images. Recent studies [14, 28, 51] have demonstrated the efficacy of masked modeling in learning effective representations from large-scale vision datasets. Our work extends this concept, discovering that masked modeling is beneficial for learning shared information across different views, with a higher masked ratio enhancing model generalization. To the best of our knowledge, this is the first attempt to develop masked modeling specifically for extracting consistent representations across multiple views.

## 3. Method

Given a multi-view dataset with  $n$  samples and  $v$  views  $\mathcal{D} = \{x^1, x^2, \dots, x^v | x^i \in \mathbb{R}^{n \times d_{x^i}}\}$ , where  $d_{x^i}$  is the dimensionality of  $i$ -th view. Our objective is to learn high-quality view-consistent representations and view-specific representations from  $\mathcal{D}$ . To simplify the notation throughout this paper, the subscript denotes  $j^{th}$  data sample, and

the superscript  $i$  indexes the view.

### 3.1. Overall Architecture

As shown in Fig. 2, our approach employs the two-stage pipeline to obtain high-quality consistency and specificity in the multi-view data. In the first stage, we use the random masking technique to process multi-view data and obtain masked multi-view samples. Then, we use a consistent encoder  $E_c$ , to extract view-consistent representations  $\mathbf{c}$  from the visible blocks. These view-consistent representations then serve as inputs for view-specific decoders, generating reconstructed views. We refer to this process as *Masked Cross-view Prediction*. In the second stage, we leverage a series of view-specific encoders  $\{E_s^i\}_{i=1}^v$  to extract view-specific representations  $\{\mathbf{s}^i\}_{i=1}^v$ . After that, we use a disentangling module to minimize the mutual information (MI) between the view-consistent representations  $\mathbf{c}$  and the view-specific representations  $\{\mathbf{s}^i\}_{i=1}^v$ . Note that, for interpretability of the representations, we require that the consistency and specificity obey the standard normal distributions  $\mathcal{N}(\mathbf{0}, \mathbf{I})$ . Below, we depict the details of mining consistency and specificity, respectively.

### 3.2. Masked Cross-view Prediction Consistency

Multi-view consistency encapsulates the shared information across different views. Building on this foundational concept, we deduce a critical implication: an effective view-consistent representation has the capacity to predict diverse views of the same object concurrently, devoid of supplementary information. Consequently, our objective is to obtain a representation that seamlessly encapsulates distinct views. However, the presence of view-specific content inherent to each view can hinder the extraction of shared information. To address this challenge, we employ the masked technique, randomly obscuring a portion of the original signals from multiple views, as illustrated in Fig. 2. Subsequently, we utilize the visible blocks as input for the consistent encoder, facilitating the learning of view-consistent representations to predict each view. The intuition behind this approach is that random masking diminishes view-specific information, compelling the consistent encoder to discern consistency in the visible segments. Moreover, the masked prediction can be viewed as a denoising process, thereby enhancing the robustness of view-consistent representations. In practical applications, the choice of the consistency encoder determines its approach. If based on Transformers [44], it can directly handle visible blocks using patch techniques. Conversely, if employing convolutional networks [13], it can directly process the masked data.

To enhance the interpretability of consistency, we assume that the prior for view-consistent representations is standard Gaussian distribution, i.e.,  $p(\mathbf{c}) \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ . There-

fore, the approximate posterior  $q_\phi(\mathbf{c}|\{x^i\}_{i=1}^v)$  can be expressed as:

$$q_\phi(\mathbf{c}|\{x^i\}_{i=1}^v) = \prod_{i=1}^v p_\phi(\mathbf{c}|x^i), \quad (1)$$

where  $\phi$  is the trainable parameters of consistent encoder  $E_c(\cdot)$ . However, Eq. (1) is intractable to optimize. To address it, we can use the reparameterization trick [24] to re-write the equality:

$$q_\phi(\mathbf{c}|\{x^i\}_{i=1}^v) = \mathcal{N}(\mu^c, (\sigma^c)^2) = \mu^c + \sigma^c \epsilon^c, \quad (2)$$

where  $\epsilon^c \sim \mathcal{N}(0, 1)$ .  $\mu^c$  and  $\sigma^c$  are parameterized with neural networks. For convenience, we can directly employ vanilla variational autoencoders (VAEs) [24] to construct MCP. To further ensure that the consistent encoder can learn multi-view consistency, we adopt a series of decoders to store view-specific content. Overall, we can derive the evidence lower bound (ELBO) of multi-view consistency:

$$\mathcal{L}_c = \mathbb{E}_{q(\mathbf{c}|\{\hat{x}^i\}_{i=1}^v)} [\log p(\{x^i\}_{i=1}^v|\mathbf{c})] - \mathbf{KL}(q_\phi(\mathbf{c}|\{\hat{x}^i\}_{i=1}^v)||p(\mathbf{c})), \quad (3)$$

where  $\{\hat{x}^i\}_{i=1}^v$  is the masked version of  $\{x^i\}_{i=1}^v$ .  $\mathbf{KL}(\cdot)$  is the KL divergence. The purpose of Eq. (3) is to learn effective view-consistent representations by minimizing the gap between masked data and original data while making the posterior of consistency  $q(\mathbf{c}|\{x^i\}_{i=1}^v)$  try to match its prior  $p(\mathbf{c})$ . During the inference stage, we discard the view-specific decoders and solely utilize the consistent encoder  $E_c(\cdot)$  to obtain the consistent representations, i.e.,  $\mathbf{c} = E_c(\{x^i\}_{i=1}^v)$ . The consistent encoder can accept both masked inputs and original inputs in the inference process.

### 3.3. Distilled Disentangling Specificity

View-specific content typically encompasses specific semantic details within a view, including angles, distances, and positional information. In general, a view results from the amalgamation of both view-specific and view-consistent information through intricate higher-order interactions. However, deducing these underlying higher-order interactions from observational data poses a formidable challenge. Consequently, the objective of this section is to disentangle corresponding view-specific information from diverse views in an environment characterized by unknown interactions.

We having obtained view-consistent representations in the previous section, a natural idea is to leverage this known factor to explore the remaining information in multi-views. We denote this approach as *Distilled Disentangling* (DD). Specifically, we employ the consistent encoder trained in the previous stage to generate the view-consistent representations. Then, we employ a series of view-specific



encoders  $\{E_s^i(\cdot)\}_{i=1}^v$  to obtain coarse view-specific representations for each view. Following this, we employ the DD module to obtain fine-grained view-specific representations  $\{s^i\}_{i=1}^v$ , by minimizing the upper bound of mutual information between coarse view-specific representations and view-consistent representations. Lastly, we concatenate the consistent representations and view-specific representations, i.e.,  $z^i = [\mathbf{c}, s^i]$ , serving as the input to the view-specific decoder  $D_s^i(\cdot)$ . The intuition behind this strategy is that, once view-specific representations are obtained through the DD module, we can employ the decoders to simulate higher-order interactions. This implies that the higher the reconstruction quality, the higher the quality of the disentangled representations we obtain.

Based on the previous description, we also assume that the prior of the specificity variable is a standard normal distribution, i.e.,  $p(s^i) \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ . Similar to the process of modeling consistency, we have  $g_{\theta^i}(s^i|x^i)$  as the posterior for specificity, where  $\theta^i$  is the trainable parameters of the view-specific encoder  $E_s^i(\cdot)$ . Similar to Eq. (2), we also have the equality:  $g_{\theta^i}(s^i|x^i) = \mu^i + \sigma^i \epsilon^i$ . To obtain disentangled view-specific representations, we need to minimize the mutual information upper bound between  $s^i$  and  $\mathbf{c}$ :

$$\begin{aligned} I(s^i; \mathbf{c}) &\equiv \mathbb{E}_{(s^i, \mathbf{c}) \sim p_{s^i, \mathbf{c}}(\cdot)} \left[ \log \frac{p(s^i|\mathbf{c})}{p(s^i)} \right] \\ &= \mathbb{E}_{(s^i, \mathbf{c}) \sim p_{s^i, \mathbf{c}}(\cdot)} \left[ \log \frac{p(s^i|\mathbf{c})}{g(s^i)} \right] - \mathbf{KL}[p(s^i) \| g(s^i)] \\ &\leq \mathbb{E}_{(s^i, \mathbf{c}) \sim p_{s^i, \mathbf{c}}(\cdot)} \left[ \log \frac{p(s^i|\mathbf{c})}{g(s^i)} \right] \\ &= \mathbf{KL}[p(s^i|\mathbf{c}) \| g(s^i)], \end{aligned} \quad (4)$$

where  $g(\cdot)$  is a variational marginal approximation. This inequality has two limitations: i) it necessitates non-negativity of the KL-divergence, and ii) learning a good marginal approximation  $g(s^i)$  to match  $p(s^i)$  is a challenge. To address it, we introduce the CLUB estimator [8] which use a variational distribution  $q_{\theta^i}(s^i|\mathbf{c})$  to approximate  $p(s^i|\mathbf{c})$ , more details given by [8]:

$$\begin{aligned} \mathcal{L}_d^i = I_{\text{CLUB}}(s^i; \mathbf{c}) &:= \mathbb{E}_{(s^i, \mathbf{c}) \sim p_{s^i, \mathbf{c}}(\cdot)} [\log q_{\theta^i}(s^i|\mathbf{c})] \\ &\quad - \mathbb{E}_{\mathbf{c} \sim p_{\mathbf{c}}(\cdot)} \mathbb{E}_{s^i \sim p_{s^i}(\cdot)} [\log q_{\theta^i}(s^i|\mathbf{c})] \quad (5) \\ &\geq I(s^i; \mathbf{c}). \end{aligned}$$

The benefit of Eq. (5) is that we can obtain disentangled representations indirectly by minimizing  $\mathcal{L}_d$ , which equally minimize the mutual information upper bound between  $s^i$  and  $\mathbf{c}$ . Additionally, we can conveniently estimate  $q_{\theta^i}(s^i|\mathbf{c})$  using neural networks.

To prevent the model obtaining trivial solutions during the disentangling process, we concatenate all the disentan-

gled representations into a new vector  $z^i = [s^i, \mathbf{c}]$ , and then input  $z^i$  into the corresponding view-specific decoder  $D_s^i(\cdot)$  to obtain the reconstruction data, i.e.,  $x'^i = D_s^i(z^i)$ . The loss function for this process is as follows:

$$\mathcal{L}_r^i = \mathbb{E}_{\mathbf{z}^i \sim q_{\theta^i}(\mathbf{z}^i|x^i)} [\log p(x^i|\mathbf{z}^i)] - \mathbf{KL}[q_{\theta^i}(\mathbf{z}^i|x^i) \| p(\mathbf{z}^i)]. \quad (6)$$

Eventually, our loss function of the second stage contains two parts:

$$\mathcal{L}_s = \frac{1}{v} \sum_{i=1}^v \mathcal{L}_d^i + \mathcal{L}_r^i \quad (7)$$

where the first term is optimized to learn disentangled view-specific representations. The second term is optimized to maintain the reconstruction quality. The pseudo-code of MRDD is given in the supplement.

## 4. Experiments

### 4.1. Dataset

We evaluate the proposed method and other competitive methods using five multi-view datasets. There are: **(a) E-MNIST** [30], which is a well-known benchmark dataset consisting of 70,000 grayscale digit images (0-9) with  $32 \times 32$  pixels. The views contain the original digits and the edge-detected version, respectively; **(b) E-FMNIST** [52], which is a fashion dataset consisting of  $32 \times 32$  grayscale images of clothing items. We synthesize the second view by running the same edge detector used to create E-MNIST; **(c) COIL-20** [33] and **(d) COIL-100** [33] which depicts from different angles containing grayscale images of 20 items and RGB images of 100 items, respectively. We create a three-view dataset by randomly grouping the images for an item into groups of three; **(f) Office-31** [38], which consists of objects commonly encountered in office settings, such as keyboards, file cabinets, and laptops. We leverage the ColorJitter method to construct a three-view dataset. We report the dataset description in the supplement.

### 4.2. Baseline Models

We evaluate two versions of the proposed method: MRDD- $\mathbf{c}$  represents the use of only view-consistent representations  $\mathbf{c}$ , while MRDD- $\mathbf{cs}$  denotes the use of  $\mathbf{c}$  and the first view-specific representations  $s^1$ . We set up three categories of baseline models for comparison with MRDD, including **(i) single-view disentangling methods:** Joint-VAE [10] and  $\beta$ -VAE [15]; **(ii) model-based multi-view disentangling methods:** MFLVC [57], EAMC [61], CONAN [18], CMC [41], and GCFAgg [59]; and **(iii) information-theory-based multi-view disentangling methods:** Multi-VAE [56], MIB [11], DVIB [3], and UNITER [58]. Our method falls into the third category. We

Method	E-MNIST		E-FMNIST		COIL-20		COIL-100		Office-31	
	ACC <sub>clu</sub>	NMI	ACC <sub>clu</sub>	NMI	ACC <sub>clu</sub>	NMI	ACC <sub>clu</sub>	NMI	ACC <sub>clu</sub>	NMI
Joint-VAE [10]	42.81±0.03	35.45±0.05	37.22±0.58	26.94±0.31	61.98±2.60	74.14±0.82	55.85±1.49	77.53±0.61	25.19±0.60	29.30±0.34
$\beta$ -VAE [15]	39.69±0.72	24.97±0.18	39.76±0.02	38.37±0.05	35.80±0.71	44.67±0.86	24.02±0.23	40.96±0.26	11.89±0.44	13.53±0.09
MFLVC† [57]	65.98±0.06	59.08±0.04	48.48±0.15	46.62±0.12	36.98±2.07	67.16±1.99	35.03±1.14	73.19±1.16	32.09±0.21	29.39±0.07
CONAN† [18]	50.22±0.02	44.42±0.01	48.70±0.11	41.41±0.01	55.94±0.88	63.98±0.46	54.04±1.37	74.73±0.47	13.51±0.24	17.11±0.22
CMC† [41]	64.16±1.04	59.84±0.77	50.13±0.42	46.24±0.41	58.03±1.18	71.23±1.22	57.19±0.94	78.16±1.07	34.28±0.32	29.75±0.19
EAMC† [61]	49.17±0.32	46.28±0.34	45.44±1.08	42.76±1.03	58.19±1.93	75.13±1.21	60.31±0.84	73.13±0.91	33.16±0.19	30.08±0.13
GCFAgg† [59]	67.10±0.88	61.34±0.62	43.09±0.07	40.25±0.21	55.79±1.66	75.08±1.38	45.71±1.35	70.22±1.82	31.17±0.21	28.54±0.31
Multi-VAE [56]	60.74±0.23	59.03±0.18	53.16±0.14	54.47±0.06	65.77±1.04	78.22±1.03	48.87±0.03	45.29±0.15	31.27±0.27	27.84±0.39
MIB [11]	53.16±0.43	48.06±1.03	<u>53.71±0.72</u>	52.66±0.59	53.60±0.89	<u>80.61±0.76</u>	53.27±0.13	<u>81.53±1.05</u>	33.08±0.26	29.94±0.35
DVIB [3]	47.80±0.01	33.29±0.02	30.23±0.85	43.36±0.30	57.47±1.55	68.27±0.95	36.75±0.48	25.54±0.36	29.24±0.38	28.69±0.21
UNITER [58]	65.35±0.11	59.71±0.25	53.07±0.83	51.19±0.10	<u>67.54±0.99</u>	79.60±1.41	50.12±0.44	54.77±0.33	34.40±0.02	<u>30.22±0.07</u>
MRDD-c (Ours)	70.34±0.05	63.49±0.02	50.11±0.09	53.55±0.08	63.96±0.03	73.88±0.03	<b>65.29±0.16</b>	<b>84.95±0.07</b>	30.06±0.59	28.14±0.15
MRDD-cs (Ours)	<b>75.93±0.12</b>	<b>69.00±0.41</b>	<b>58.25±0.27</b>	<b>59.93±0.39</b>	<b>69.18±0.44</b>	<b>81.52±0.14</b>	62.00±0.27	83.65±1.04	<b>37.14±0.98</b>	<b>39.26±0.53</b>
$\Delta$ SOTA	+8.83	+7.66	+4.54	+4.86	+1.64	+0.91	+4.98	+3.34	+2.86	+9.04

Table 1. **Clustering results (%) on five datasets.** **Bold** denotes the best results and underline denotes the second-best. † denotes we set the dimensionality of latent representations as 10. All results are reproduced using the officially released code.

evaluate the clustering and classification performance of all comparative models using k-means and support vector classification (SVC). Each model undergoes 10 runs, and we report their average values and variances. Notably, for single-view methods, we select the results from the best view as the evaluation outcome. In the case of multi-view methods limited to two views, we choose the optimal two views as inputs for these models, such as MIB and DVIB.

### 4.3. Implementation Details

We implement the proposed method and other non-linear comparison methods on the PyTorch 2.0.1 [34] platform, running on Ubuntu 18.04 LTS utilizing one GPU (NVIDIA GeForce RTX 2080 Ti with 12 GB of memory). For simplicity, we use convolutional networks to build the encoders and decoders<sup>2</sup>. For all experiments, we default the masked ratio of MCP to 70%, and both the dimensions of consistency and specificity are set to 10. It is worth noting that the MCP’s masked ratio is set at 70% to strike a balance between minimizing information loss and optimizing encoder computation cost. We train two stages of our model for 200 epochs, respectively. We use the Adam optimizer with a batch size of 512 for the model and then employ the cosine annealing learning rate scheduler with an initial learning rate of  $5 \times 10^{-4}$ . To ensure a fair comparison, we use the recommended optimal settings from the original paper for other baseline models. In the classification task, we split the dataset into a training set and a test set in a ratio of 80:20.

### 4.4. Evaluation Metrics

In order to evaluate clustering performance, two standard evaluation metrics are used: clustering Accuracy

<sup>2</sup>Network structure details are given in the supplement.

Method	E-MNIST		COIL-100		Office-31	
	ACC <sub>cls</sub>	F-Score	ACC <sub>cls</sub>	F-Score	ACC <sub>cls</sub>	F-Score
Random	10.00±0.08	10.00±0.08	1.10±0.15	1.08±0.13	3.17±0.15	3.13±0.14
Joint-VAE [10]	81.16±0.18	80.97±0.20	<u>86.73±0.86</u>	<u>85.94±0.80</u>	42.32±1.10	40.51±1.15
$\beta$ -VAE [15]	43.82±0.38	41.84±0.39	30.87±1.15	25.64±0.83	28.09±0.75	25.25±0.94
CONAN† [18]	61.42±0.36	59.07±0.36	62.45±1.30	57.33±0.61	47.07±1.47	44.79±1.41
CMC† [41]	93.86±0.04	93.05±0.01	80.29±0.51	79.48±1.10	51.33±0.46	50.87±0.62
Multi-VAE [56]	92.73±0.16	92.76±0.16	74.11±2.18	70.78±0.95	62.53±1.17	62.20±1.39
MIB [11]	90.81±0.30	90.03±0.49	79.01±0.04	78.63±0.09	61.48±0.43	60.79±0.39
DVIB [3]	84.14±0.19	83.98±0.24	<u>67.22±0.73</u>	66.34±1.16	58.12±0.87	56.52±0.77
UNITER [58]	94.21±0.16	94.11±0.38	84.26±1.05	84.23±1.21	69.33±1.42	67.51±0.99
MRDD-c (Ours)	93.69±0.17	93.67±0.18	86.67±1.65	84.86±1.96	53.04±0.83	48.92±1.38
MRDD-cs (Ours)	<b>98.37±0.09</b>	<b>98.36±0.09</b>	<b>91.17±1.21</b>	<b>90.06±1.15</b>	<b>73.51±0.28</b>	<b>72.71±0.44</b>
$\Delta$ SOTA	+4.16	+4.25	+4.44	+4.12	+4.18	+5.2

Table 2. **Classification results (%) on three datasets.** **Bold** denotes the best results and underline denotes the second-best. † denotes we set the dimensionality of latent representations as 10. All results are reproduced using the officially released code.

(ACC<sub>clu</sub>) and Normalized Mutual Information (NMI). Readers seeking further details on these metrics are referred to [25]. It is important to note that the validation process of clustering methods is limited to cases where ground truth labels are available. For classification, ACC<sub>cls</sub> and F-Score are used as evaluation metrics. In all cases, a high value indicates better performance<sup>3</sup>.

### 4.5. Comparison Results and Analysis

We evaluate the performance of all baseline models on five datasets for both clustering and classification tasks<sup>4</sup>, as shown in Tables 1 and 2. These results indicate that, under the same experimental settings, the representations extracted by our method significantly improve the perfor-

<sup>3</sup>The details of the evaluation metrics are given in the supplementary materials.

<sup>4</sup>More results are given in the supplement.

Method	E-MNIST	E-FMNIST	COIL-20	COIL-100	Office-31
MRDD-cs (baseline)	<b>98.37</b>	<b>88.78</b>	<b>95.97</b>	<b>91.17</b>	<b>73.51</b>
only stage I	93.69	82.51	88.18	86.67	53.04
only stage II	62.13	58.64	69.65	61.28	40.68
w/o MCP	93.12	83.11	89.48	86.93	61.45
w/o $\mathcal{L}_d$	94.57	83.07	90.72	85.37	58.92
w/o $\mathcal{L}_r$	90.63	76.88	86.90	88.96	50.04

Table 3. **Components ablation study.** Classification accuracy scores (%) on five dataset. Best results are in **bold**.

mance of downstream tasks. At the same time, we find that MRDD-c exhibits superior performance in both  $ACC_{clu}$  and NMI metrics compared to MRDD-cs. This suggests that naively concatenating features may lead to a degradation in the performance of the representation. In addition, the results on COIL-20 and COIL-100 indicate that Joint-VAE outperforms the majority of multi-view methods, such as MFLVC, CMC, and GCFAgg, in terms of evaluation metrics. This outcome highlights that the quality of views can impact the quality of representations, consequently leading to a degradation in the performance of fused representations.

#### 4.6. Ablation Study

**Components study.** We conducted ablation experiments to evaluate the effects of different components in the proposed method MRDD-cs, as shown in Table 3. In the case of using only Stage I and Stage II respectively indicate that our framework is superior to any of the two independent components. Moreover, the results of using only Stage II support our hypothesis that it is difficult to extract high-quality multi-view representations in an unsupervised environment without additional information. On the other hand, the ablation results show that removing MCP will degrade the performance of our method. Notably, the ablation experiments of  $\mathcal{L}_d$  and  $\mathcal{L}_r$  demonstrate that removing  $\mathcal{L}_r$  has a more significant negative impact on our method.

**Masked ratio and strategies.** In this study, we conduct a series of ablation studies focusing on the MCP component, applying various masked ratios and masking strategies. The results of Fig. 3 show a notable trend: an increase in the masked ratio correlates positively with the model’s classification performance. It is observed that, for the majority of the datasets examined, an optimal performance is attained at a masked ratio of 70%. This observation aligns well with the findings reported in prior studies. Interestingly, specific datasets, notably COIL-20 and Office-31, exhibited a peak in classification accuracy at a high masked ratio of 80%. Furthermore, it was observed that even at a masked ratio of 90%, the decrement in performance was not pronounced. This resilience in performance at high masked ratios can be postulated to stem from the multi-view context employed, wherein the model retains the capability to assimilate valuable representations from the visible blocks




Figure 3. **Masked ratio.** Classification accuracy scores (%) for masked ratios range from 0% to 90% on five datasets.

Strategy	E-MNIST	E-FMNIST	COIL-20	COIL-100	Office-31
Random	<b>98.37</b>	<b>88.78</b>	<b>95.97</b>	<b>91.17</b>	<b>73.51</b>
Block-wise [12]	96.90	87.24	94.55	89.79	70.14
Grid-wise [14]	97.15	87.62	93.19	90.03	71.43

Table 4. **Masked strategy.** Classification results (%) of three masked strategies (e.g., random, block-wise, and grid-wise) with the masked ratio of 70% on five datasets. Best results are in **bold**.

of alternate views.

Furthermore, our research included an exploration of three distinct masking strategies—random, block-wise, and grid-wise—to evaluate their respective impacts on the model’s performance. By default, our methodology utilized the random strategy. The comparative analysis of these strategies, as depicted in Table 4, indicates that the differences in performance metrics among the three strategies are relatively marginal. Notably, the random strategy outperformed the others across all datasets. The performance of the block-wise strategy was observed to be inferior compared to the other two. We conjecture that this underperformance may be attributed to the block-wise strategy’s inherent design, which involves the removal of large blocks, potentially leading to the concurrent loss of shared information.

**The dimensionality of consistency and specificity.** We study the effects of view-consistent and view-specific representations extracted by our method in different dimensions<sup>5</sup>. We set the range of view-consistent representation dimensions to  $\{5, 10, 15, 20\}$  and the view-specific representation dimension to  $\{5, 10, 15, 20, 40\}$ . Fig. 4 shows that when the consistent representation dimension is fixed, the clustering performance is positively correlated with the view-specific representation dimension. Our method achieved the best clustering performance on the E-MNIST and E-FMNIST datasets when the consistent representation dimension is 15 and the specific representation dimension is 40. On the other hand, Fig. 4 indicates that a large dimension of consistent representations can lead to performance degradation. For example, when the consistent representation dimension is set to 20, the clustering results were inferior to the consistent representation dimension of 15.

<sup>5</sup>More results are given in the supplement.

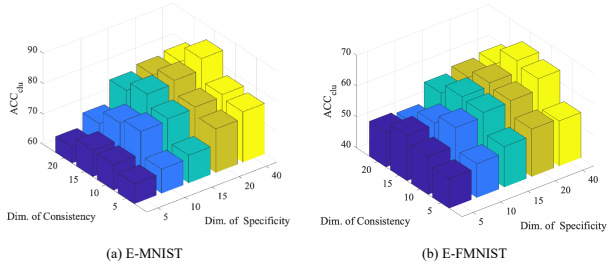


Figure 4. The clustering results (%) of the different dimensions of consistency and specificity on the E-MNIST and E-FMNIST datasets. The x-axis represents the consistency dimension, the y-axis represents the specificity dimension, and the z-axis represents the clustering accuracy.

Type	E-MNIST	E-FMNIST	COIL-20	COIL-100	Office-31
c	70.34	50.11	63.96	<b>65.29</b>	30.06
$s^1$	55.93	<b>59.73</b>	63.14	60.29	34.42
$s^2$	43.32	59.40	63.44	59.92	32.51
$s^3$	-	-	63.19	61.46	33.80
c + $s^1$	<b>75.93</b>	58.25	<b>69.18</b>	62.00	37.14
concat.	73.16	58.10	63.26	60.90	<b>40.84</b>

Table 5. The clustering results (%) of different types of representations on five datasets, where c represents view-consistent representations,  $s^i$  represents the  $i$ -th view-specific representations, and *concat.* represents concatenating all representations. Best results are in **bold**.

**The contributions of consistency and specificity.** We evaluate the clustering results of view-consistent and view-specific representations on five datasets. From Table 5, we note that the manner in which features are combined exerts a pronounced effect on the performance of downstream tasks. This is attributable to the fact that we have generated additional views by applying data augmentation to a single view, resulting in the first view containing a greater quantity of information. Consequently, the use of just the c +  $s^1$  feature combination can achieve superior performance. On the other hand, we also observe that a simple concatenation of all features introduces redundancy, which in turn diminishes the performance of downstream models. These experimental findings suggest that a singular approach to feature fusion is unlikely to have a universally positive impact on all tasks.

#### 4.7. Visualization

We visualize the representations of MRDD-c and MRDD-cs on the E-MNIST and E-FMNIST datasets<sup>6</sup>. Fig. 5 indicates that view-consistent representations can distinguish different samples at a coarse level. However, after incorporating view-specific representations, the discriminative ability of the representations is enhanced, especially evident in the E-MNIST dataset.

<sup>6</sup>More visualization results are given in the supplement.

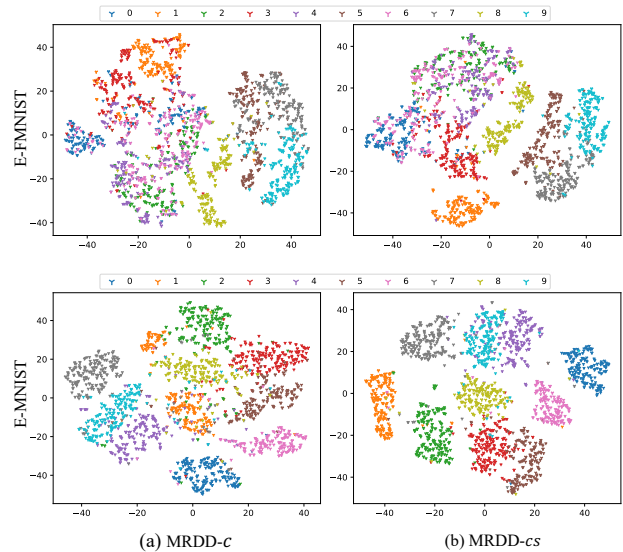


Figure 5. Visualization of the representations of MRDD-c and MRDD-cs using t-SNE [43] on the E-MNIST and E-FMNIST datasets.

## 5. Discussion and Conclusion

This work attempts to address the problem of obtaining high-quality multi-view consistent and specific representations from the perspective of representation disentangling. Extensive experimental results reveal that: i) A high masked ratio can enhance the performance of view-consistent representations while reducing computational costs. ii) Diminishing the size of the view-consistent representations in relation to view-specific representations substantially enhances the efficacy of their integrated representations. iii) The role of view-specific representations in predictive tasks is significant, and the mere aggregation of all representations might lead to suboptimal performance.

Although the proposed method is adept at disentangling consistency and specificity, the efficacy of this decoupling is contingent upon the quality of the consistent representations. On the other hand, while the MCP enhances performance and efficiency, it also increases the instability of training. In complex scenarios, the consistent representations obtained by MCP make it challenging to reconstruct the original data. Our future work will consider how to reduce the dependency between the two stages and improve the stability of the MCP module.

**Acknowledgment.** This work was supported by the Beijing Natural Science Foundation (No. 4234086); the Natural Science Foundation of China (No. 62192782); Guangdong Natural Science Funds for Distinguished Young Scholar (No. 2023B1515020097); Singapore MOE Tier 1 Funds (MSS23C002); and the NRF Singapore under the AI Singapore Programme (No. AISG3-GV-2023-011).



## References

- [1] Galen Andrew, Raman Arora, Jeff Bilmes, and Karen Livescu. Deep canonical correlation analysis. In *ICML*, pages 1247–1255, 2013. 3
- [2] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A framework for self-supervised learning of speech representations. *NeurIPS*, 33:12449–12460, 2020. 3
- [3] Feng Bao. Disentangled variational information bottleneck for multiview representation learning. In *CICAI*, pages 91–102, 2021. 1, 3, 5, 6
- [4] Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeshwar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and Devon Hjelm. Mutual information neural estimation. In *ICML*, pages 531–540, 2018. 1
- [5] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE TPAMI*, 35(8):1798–1828, 2013. 3
- [6] Guoqing Chao and Shiliang Sun. Consensus and complementarity based maximum entropy discrimination for multi-view classification. *Information Sciences*, 367:296–310, 2016. 3
- [7] Jiawei Chen and Chiu Man Ho. Mm-vit: Multi-modal video transformer for compressed video action recognition. In *CVPR*, pages 1910–1921, 2022. 1
- [8] Pengyu Cheng, Weituo Hao, Shuyang Dai, Jiachang Liu, Zhe Gan, and Lawrence Carin. Club: A contrastive log-ratio upper bound of mutual information. In *ICML*, pages 1779–1788, 2020. 5
- [9] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *CVPR*, pages 886–893, 2005. 1, 3
- [10] Emilien Dupont. Learning disentangled joint continuous and discrete representations. *NeurIPS*, 31, 2018. 5, 6
- [11] Marco Federici, Anjan Dutta, Patrick Forré, Nate Kushman, and Zeynep Akata. Learning robust representations via multi-view information bottleneck. *ICLR*, 2020. 1, 3, 5, 6
- [12] Songhao Piao Furu Wei Hangbo Bao, Li Dong. Beit: BERT pre-training of image transformers. In *ICLR*, 2022. 7
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 4
- [14] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, pages 16000–16009, 2022. 3, 7
- [15] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *ICLR*, 2016. 5, 6
- [16] Zhenyu Huang, Joey Tianyi Zhou, Hongyuan Zhu, Changqing Zhang, Jiancheng Lv, and Xi Peng. Deep spectral representation learning from multi-view data. *IEEE TIP*, 30: 5352–5362, 2021. 3
- [17] Qin Jin, Jia Chen, Shizhe Chen, Yifan Xiong, and Alexander Hauptmann. Describing videos using multi-modal fusion. In *ACM MM*, pages 1087–1091, 2016. 1
- [18] Guanzhou Ke, Zhiyong Hong, Zhiqiang Zeng, Zeyi Liu, Yangjie Sun, and Yannan Xie. Conan: contrastive fusion networks for multi-view clustering. In *IEEE International Conference on Big Data*, pages 653–660, 2021. 1, 2, 3, 5, 6
- [19] Guanzhou Ke, Zhiyong Hong, Wenhua Yu, Xin Zhang, and Zeyi Liu. Efficient multi-view clustering networks. *Applied Intelligence*, 52(13):14918–14934, 2022. 3
- [20] Guanzhou Ke, Yongqi Zhu, and Yang Yu. Mori-ran: Multi-view robust representation learning via hybrid contrastive fusion. In *IEEE International Conference on Data Mining Workshops*, pages 467–474, 2022. 3
- [21] Guanzhou Ke, Guoqing Chao, Xiaoli Wang, Chenyang Xu, Yongqi Zhu, and Yang Yu. A clustering-guided contrastive fusion for multi-view representation learning. *IEEE TCSVT*, 2023. 2
- [22] Guanzhou Ke, Yang Yu, Guoqing Chao, Xiaoli Wang, Chenyang Xu, and Shengfeng He. Disentangling multi-view representations beyond inductive bias. In *ACM MM*, pages 2582–2590, 2023. 1, 3
- [23] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, page 2, 2019. 3
- [24] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2014. 4
- [25] Abhishek Kumar, Piyush Rai, and Hal Daume. Co-regularized multi-view spectral clustering. *NeurIPS*, 24: 1413–1421, 2011. 6
- [26] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. ALBERT: A lite BERT for self-supervised learning of language representations. In *ICLR*, 2020. 3
- [27] Yingming Li, Ming Yang, and Zhongfei Zhang. A survey of multi-view representation learning. *TKDE*, 31(10):1863–1883, 2019. 3
- [28] Yanghao Li, Haoqi Fan, Ronghang Hu, Christoph Feichtenhofer, and Kaiming He. Scaling language-image pre-training via masking. In *CVPR*, pages 23390–23400, 2023. 3
- [29] Jialu Liu, Chi Wang, Jing Gao, and Jiawei Han. Multi-view clustering via joint nonnegative matrix factorization. In *ICDM*, pages 252–260, 2013. 3
- [30] Ming-Yu Liu and Oncel Tuzel. Coupled generative adversarial networks. *NeurIPS*, 29, 2016. 5
- [31] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019. 3
- [32] David G Lowe. Object recognition from local scale-invariant features. In *ICCV*, pages 1150–1157, 1999. 1
- [33] Sameer A Nene, Shree K Nayar, Hiroshi Murase, et al. Columbia object image library (coil-20). 1996. 5
- [34] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *NeurIPS*, 32:8026–8037, 2019. 6

- [35] Nikhil Rasiwasia, Jose Costa Pereira, Emanuele Coviello, Gabriel Doyle, Gert RG Lanckriet, Roger Levy, and Nuno Vasconcelos. A new approach to cross-modal multimedia retrieval. In *ACM MM*, pages 251–260, 2010. 1, 3
- [36] Sucheng Ren, Yong Du, Jianming Lv, Guoqiang Han, and Shengfeng He. Learning from the master: Distilling cross-modal advanced knowledge for lip reading. In *CVPR*, pages 13325–13333, 2021. 3
- [37] Sucheng Ren, Zhengqi Gao, Tianyu Hua, Zihui Xue, Yonglong Tian, Shengfeng He, and Hang Zhao. Co-advise: Cross inductive bias distillation. In *CVPR*, pages 16773–16782, 2022. 3
- [38] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *ECCV*, pages 213–226, 2010. 5
- [39] Shiliang Sun and Guoqing Chao. Multi-view maximum entropy discrimination. In *IJCAI*, pages 1706–1712, 2013. 3
- [40] Xiukun Sun, Miaomiao Cheng, Chen Min, and Liping Jing. Self-supervised deep multi-view subspace clustering. In *ACML*, pages 1001–1016, 2019. 3
- [41] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *ECCV*, pages 776–794, 2020. 3, 5, 6
- [42] Daniel J Trosten, Sigurd Lokse, Robert Jenssen, and Michael Kampffmeyer. Reconsidering representation alignment for multi-view clustering. In *CVPR*, pages 1255–1265, 2021. 2, 3
- [43] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. 8
- [44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 30, 2017. 4
- [45] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *ICML*, pages 1096–1103, 2008. 3
- [46] Weiran Wang, Raman Arora, Karen Livescu, and Jeff Bilmes. On deep multi-view representation learning. In *ICML*, pages 1083–1092, 2015. 1, 3
- [47] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training. In *CVPR*, pages 3024–3033, 2021. 3
- [48] Xin Wang, Hong Chen, Si’ao Tang, Zihao Wu, and Wenwu Zhu. Disentangled representation learning. *arXiv preprint arXiv:2211.11695*, 2022. 3
- [49] Xiaoli Wang, Yongli Wang, Guanzhou Ke, Yupeng Wang, and Xiaobin Hong. Knowledge distillation-driven semi-supervised multi-view classification. *Information Fusion*, 103:102098, 2024. 2
- [50] Yang Wang, Lin Wu, Xuemin Lin, and Junbin Gao. Multi-view spectral clustering via structured low-rank matrix factorization. *IEEE TNNLS*, 29(10):4833–4843, 2018. 3
- [51] Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan Yuille, and Christoph Feichtenhofer. Masked feature prediction for self-supervised visual pre-training. In *CVPR*, pages 14668–14678, 2022. 3
- [52] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017. 5
- [53] Yi Xie, Huaidong Zhang, Xuemiao Xu, Jianqing Zhu, and Shengfeng He. Towards a smaller student: Capacity dynamic distillation for efficient image retrieval. In *CVPR*, pages 16006–16015, 2023. 3
- [54] Yi Xie, Yihong Lin, Wenjie Cai, Xuemiao Xu, Huaidong Zhang, Yong Du, and Shengfeng He. D3still: Decoupled differential distillation for asymmetric image retrieval. In *CVPR*, 2024. 3
- [55] Cai Xu, Wei Zhao, Jinglong Zhao, Ziyu Guan, Yaming Yang, Long Chen, and Xiangyu Song. Progressive deep multi-view comprehensive representation learning. In *AAAI*, pages 10557–10565, 2023. 3
- [56] Jie Xu, Yazhou Ren, Huayi Tang, Xiaorong Pu, Xiaofeng Zhu, Ming Zeng, and Lifang He. Multi-vae: Learning disentangled view-common and view-peculiar visual representations for multi-view clustering. In *CVPR*, pages 9234–9243, 2021. 1, 2, 3, 5, 6
- [57] Jie Xu, Huayi Tang, Yazhou Ren, Liang Peng, Xiaofeng Zhu, and Lifang He. Multi-level feature learning for contrastive multi-view clustering. In *CVPR*, pages 16051–16060, 2022. 1, 2, 3, 5, 6
- [58] Jie Xu, Yazhou Ren, Xiaoshuang Shi, Heng Tao Shen, and Xiaofeng Zhu. Untie: Clustering analysis with disentanglement in multi-view information fusion. *Information Fusion*, 100:101937, 2023. 5, 6
- [59] Weiqing Yan, Yuanyang Zhang, Chenlei Lv, Chang Tang, Guanghui Yue, Liang Liao, and Weisi Lin. Gcfagg: Global and cross-view feature aggregation for multi-view clustering. In *CVPR*, pages 19863–19872, 2023. 5, 6
- [60] Changqing Zhang, Yeqing Liu, and Huazhu Fu. Ae2-nets: Autoencoder in autoencoder networks. In *CVPR*, pages 2577–2585, 2019. 3
- [61] Runwu Zhou and Yi-Dong Shen. End-to-end adversarial-attention network for multi-modal clustering. In *CVPR*, pages 14619–14628, 2020. 2, 3, 5, 6
- [62] Yin Zhou, Pei Sun, Yu Zhang, Dragomir Anguelov, Jiyang Gao, Tom Ouyang, James Guo, Jiquan Ngiam, and Vijay Vasudevan. End-to-end multi-view fusion for 3d object detection in lidar point clouds. In *CoRL*, pages 923–932, 2020. 1