

Multi-View Factorizing and Disentangling: A Novel Framework for Incomplete Multi-View Multi-Label Classification

Wulin Xie
Guizhou University
China

xwl1085930920@gmail.com

Lian Zhao
Guizhou University
China

2816920140@qq.com

Jiang Long
Guizhou University
China

longj_2000@163.com

Xiaohuan Lu
Guizhou University
China

xhlu3@gzu.edu.cn

Bingyan Nie*
Guizhou University
China

bingyannie966@gmail.com

Abstract

Multi-view multi-label classification (MvMLC) has recently garnered significant research attention due to its wide range of real-world applications. However, incompleteness in views and labels is a common challenge, often resulting from data collection oversights and uncertainties in manual annotation. Furthermore, the task of learning robust multi-view representations that are both view-consistent and view-specific from diverse views still a challenge problem in MvMLC. To address these issues, we propose a novel framework for incomplete multi-view multi-label classification (iMvMLC). Our method factorizes multi-view representations into two independent sets of factors: view-consistent and view-specific, and we correspondingly design a graph disentangling loss to fully reduce redundancy between these representations. Additionally, our framework innovatively decomposes consistent representation learning into three key sub-objectives: (i) how to extract view-shared information across different views, (ii) how to eliminate intra-view redundancy in consistent representations, and (iii) how to preserve task-relevant information. To this end, we design a robust task-relevant consistency learning module that collaboratively learns high-quality consistent representations, leveraging a masked cross-view prediction (MCP) strategy and information theory. Notably, all modules in our framework are developed to function effectively under conditions of incomplete views and labels, making our method adaptable to various multi-view and multi-label datasets. Extensive experiments on five datasets demonstrate that our method outperforms other leading approaches.

*Corresponding author: bingyannie966@gmail.com

1. Introduction

Nowadays, multi-view data, collected from diverse sources such as SIFT, Gist, and HSV or obtained from various feature extractors, has garnered significant research attention due to its potential to provide richer and more diverse descriptions of samples [10] [32] [28]. This leads to an intriguing research question: how can multi-view features be effectively integrated to extract meaningful information, many multi-view learning methods have been proposed to address this question [24] [6]. Concurrently, multi-label classification has long been a crucial area in pattern recognition [25] [7] as single-label data often fail to capture the complexities of real-world scenarios. For instance, an image might be described by multiple elements such as “neon signs”, “sidewalks”, and “passing cars”, each contributing to a comprehensive understanding of the scene. Consequently, the field of multi-view learning now faces a new composite challenge: multi-view multi-label classification (MvMLC).

However, most MvMLC methods assume that all views and labels are complete, which does not consistently align with real-world scenarios. For example, if an online article contains only text and images, the option to view a video will not be available. Additionally, manual annotation may overlook certain tags due to errors or budget constraints, reducing the effectiveness of multi-label supervision. This challenge highlights the necessity of incomplete multi-view multi-label classification (iMvMLC), which focuses on scenarios where some views or labels are missing. To address this problem, several related works have been proposed in recent years. For example, Li et al. proposed NAIM3L that leveraged consensus and structural information among labels to ease the incompleteness problem. [9] Liu et al. proposed DICNet, a deep instance-level contrastive network that sings

stacked autoencoders for view-specific feature extraction and an instance-level contrastive learning scheme to enhance consensus representation. [13] Liu et al. introduced LMVCAT which leverages a transformer-based network to capture the relationship between different views and labels. [14] However, most existing iMvMLC methods rely on end-to-end frameworks to simultaneously learn consistent and view-specific representations. This approach not only creates a Min-Max game that risks converging on sub-optimal solutions in the absence of supplementary information but also significantly increases computational demands as the number of views grows [6]. Additionally, coupling view-specific and view-shared information within the same feature space inevitably leads to redundancy between the two. To ensure simultaneous learning of these different types of information during representation learning, models inevitably discarded some specific semantic details during the encoding process, resulting in the loss of some view-specific information.

To address these challenges, we propose a multi-view factorizing and disentangling framework (MVFD) for the iMvMLC task. Unlike existing approaches that utilize an end-to-end framework to concurrently learn consistent and view-specific representations, our method factorizes multi-view representations into view-consistent and view-specific factors and employs a two-stage multi-view learning paradigm to learn each separately. Furthermore, to obtain pure view-consistent representation for downstream task, we decompose consistent representation learning into three objectives: extracting view-shared information, eliminating redundant intra-view information, and preserving task-relevant information. We correspondingly design a collaborative consistency learning module based on a masked cross-view prediction (MCP) strategy and information theory. Finally, we design a graph disentangling loss to fully reduce redundancy between view-consistent and view-specific representations.

In summary, the contributions of our method can be highlighted as follows:

- We propose a novel framework for iMvMLC that factorizes multi-view representations into view-consistent and view-specific factors and employs a two-stage learning paradigm to learn them separately. Additionally, our framework is compatible with arbitrary incomplete multi-view and weak multi-label data.
- Our framework innovatively decomposes the consistency learning task into three sub-objects and combines the masked cross-view prediction (MCP) strategy and information theory for collaborative optimization.
- We design a graph disentangling loss aimed at minimizing redundant information between view-consistent and view-specific representations, which can also be adapted to current multi-view representation learning

methods. Extensive experiments on five datasets validate that our method outperforms other state-of-the-art approaches for the iMvMLC task.

2. Related Works

2.1. Multi-View Multi-Label Classification

In recent years, many works for multi-view multi-label classification (MvMLC) tasks have been proposed [15] [38] [36]. Sun et al. introduce the Latent Conditional Bernoulli Mixture (LCBM), a multi-view probabilistic model for multi-label classification that enhances label dependency modeling and generalization through a novel integration of Bernoulli mixture models and a Gaussian Mixture Variational Autoencoder (GMVAE) within a variational inference framework [22]. Zhang et al. proposed LSA-MML, which utilizes matrix factorization and the Hilbert-Schmidt Independence Criterion to align and enhance the representation of multiple views. [35]. In addition, neural networks (DNN) have also been introduced in the MvML task. For example, Wu et al. proposed SIMM that combines confusion adversarial learning and the orthogonal constraint to effectively integrate and optimize both shared and view-specific information [30]. Zhao et al. proposed CDMM that addresses the challenge of maintaining consistency and diversity among multiple views by using separate classifiers for each view, integrating the Hilbert-Schmidt Independence Criterion for diversity, and incorporating label correlation and view contribution factors to enhance prediction accuracy. [37]

2.2. Incomplete Multi-View Multi-Label Classification

However, existing MvMLC methods ignore the incompleteness problem in MvMLC and can't handle missing views or missing labels. To this end, many methods for iMvMLC have been developed that can handle missing views and missing labels simultaneously [17, 18, 21, 31]. For instance, Li et al. proposed NAIM3L that effectively handles missing labels, and incomplete and non-aligned views with minimal assumptions, employing a single hyper-parameter and a high-rank modeling approach for global label structures [9]. Tan et al. proposed iMVWL which simultaneously addresses incomplete views and weak labels by learning a shared subspace that captures cross-view relationships and local label correlations [23]. Recently, DNN has proven its huge potential in iMvMLC tasks. Liu et al. proposed DICNet that introduces a DNN-based framework for iMvMLC that effectively handles incomplete data through view-specific autoencoders [13]. Liu et al. proposed MTD that enhances consistency and uniqueness in view representations through a two-channel encoder system and utilizes a weak label-guided graph regularization to improve feature decoupling and preserve geometric structures in embeddings [12].

3. Method

In this section, we introduce our method in detail from the following three aspects, namely factorized consistent representation Learning, view-specific representation disentangling, and multi-label classification. For the convenience of description, we first introduce the formal problem definition and common notations.

3.1. Problem definition

In order to clearly describe the iMvMLC task, we define original multi-view data as $\{\mathbf{X}^{(v)} \in \mathbb{R}^{n \times d_v}\}_{v=1}^m$ where m , n , d_v denotes the number of views, sample, and the original dimension of v -th view feature, respectively. $\mathbf{Y} \in \{0, 1\}^{n \times c}$ denotes the label matrix with c categories and $\mathbf{Y}_{i,j} = 1$ denotes that the i -th sample belongs to the j -th category. Furthermore, we introduce missing view indicator $\mathbf{W} \in \{0, 1\}^{n \times m}$ and missing label indicator $\mathbf{G} \in \{0, 1\}^{n \times m}$, respectively, where $\mathbf{W}_{i,j} = 1$ means the j -th view of i -th sample is available, otherwise $\mathbf{W}_{i,j} = 0$. Similarly, $\mathbf{G}_{i,j} = 1$ represents i -th sample's j -th label is known, otherwise $\mathbf{G}_{i,j} = 0$. Additionally, we filled the missing data in $\{\mathbf{X}^{(v)} \in \mathbb{R}^{n \times d_v}\}_{v=1}^m$ and $\mathbf{Y} \in \{0, 1\}^{n \times c}$ with '0' in the data pre-processing step. The target of our method is to train a network that can accurately predict the categories of incomplete multi-view data.

3.2. Factorized Consistent Representation Learning

View-shared information extracting: As we know, the desired latent consistent representation should encode the shared information from observed views. As discussed in [34] and [6], an effective view-consistent representation can concurrently reconstruct diverse views of the same object if each view is conditionally independent given the shared multi-view representation. Based on this, we define the desired multi-view consistent representation as follows:

Proposition 1 *A multi-view consistent representation c contains the shared information across different views if each observation, i.e., $x^{(v)}$ from $x^{(1)}, \dots, x^{(m)}$, can be reconstructed from a mapping $f^{(v)}(\cdot)$, i.e., $x^{(v)} = f^{(v)}(c)$.*

where c denotes the consistent representation of multiple views. Additionally, since the original multi-view data contains redundant view-specific information that may interfere with consistency learning, we are inspired by masked autoencoders (MAE) [4], which randomly mask image patches. As shown in Fig. 1, in the first stage, we apply a random masking technique to the input multi-view vector data. Specifically, we first construct a matrix $\mathbf{M}^{(v)} \in \{0, 1\}^{n \times d_v}$ for any view v , where all elements are initially set to 1. For each row $m_{i,:} \in \mathbf{M}^{(v)}$, where $i \in \{1, 2, \dots, n\}$, we randomly select

an initial point $p_i^{(v)} \in [0, d_v - l^{(v)}]$ and set the elements to 0 from position $p_i^{(v)}$ to $p_i^{(v)} + l^{(v)}$, where $l^{(v)}$ denotes the mask length in view v . Specifically, $l^{(v)} = d_v \cdot \delta$, where δ is a parameter representing the mask ratio. Finally, we obtain the weight matrix $\bar{\mathbf{M}}^{(v)}$, and use the weight matrices $\{\bar{\mathbf{M}}^{(v)}\}_{v=1}^m$ to mask the original multi-view data as follows:

$$\{\bar{\mathbf{X}}^{(v)}\}_{v=1}^m = \{\mathbf{X}^{(v)} \odot \bar{\mathbf{M}}^{(v)}\}_{v=1}^m \quad (1)$$

where \odot denotes Hadamard product and $\{\bar{\mathbf{X}}^{(v)}\}_{v=1}^m$ is final masked multi-view data. Then we use a set of multi-layer perceptrons (MLP) as consistent encoders to extract shared information from multiple views: $\{E_c^{(v)} : \bar{\mathbf{X}}^{(v)} \in \mathbb{R}^{n \times d_v} \rightarrow \mathbf{C}^{(v)} \in \mathbb{R}^{n \times d_e}\}_{v=1}^m$ where $E_c^{(v)}$ means the consistent encoder of view v , $\mathbf{C}^{(v)} \in \mathbb{R}^{n \times d_e}$ denotes the v -th view's embedding consistent feature and d_e represents the embedding dimension. Then we obtain comprehensive consistent representation $\{\hat{\mathbf{C}}_{i,:}\}_{i=1}^n$ as follows:

$$\hat{\mathbf{C}}_{i,:} = \sum_{v=1}^m \frac{\mathbf{C}^{(v)} \mathbf{W}_{i,v}}{\sum_v \mathbf{W}_{i,v}} \quad (2)$$

where $\mathbf{W} \in \mathbb{R}^{n \times m}$ means the missing view indicator matrix. To satisfy the condition in proposition 1, we utilize a set of MLPs to construct view-specific decoders as mapping functions and put consistent embedding features as input to reconstruct raw multi-view data as follows: $\{D_c^v : \hat{\mathbf{C}} \in \mathbb{R}^{n \times d_e} \rightarrow \mathbf{X}'^{(v)} \in \mathbb{R}^{n \times d_v}\}_{v=1}^m$ where $\mathbf{X}'^{(v)}$ denotes the reconstructed feature of view v . Then we design masked consistent prediction loss as follows:

$$\ell_{cp} = \frac{1}{m} \sum_{v=1}^m \ell_{cp}^{(v)} = \frac{1}{m} \sum_{v=1}^m \left(\frac{1}{d_v} \sum_{i=1}^n \left\| \mathbf{X}'_i^{(v)} - \mathbf{X}_i^{(v)} \right\|_2^2 W_{i,v} \right) \quad (3)$$

where $W_{i,v}$ denotes the availability of the v -th view in i -th sample. By minimizing Eq.(3), $\hat{\mathbf{C}}$ can encode shared information from available views, and different samples (regardless of their missing patterns) are associated with representations in a common space.

Redundant intra-view information eliminating: While minimizing Eq.(3) can extract consistency among different views from high-level features, the reconstruction objective inevitably results in the same features to maintain the view-private information for individual view [33], which hinders the extraction of consistent information. To this end, we introduce contrastive learning into semantic space to eliminate view-specific information within high-level consistent features. Specifically, we first obtain predictions of all views through a shared classifier as follows: $\{\Psi_c : \mathbf{C}^{(v)} \in \mathbb{R}^{n \times d_e} \rightarrow \hat{\mathbf{P}}^{(v)} \in \mathbb{R}^{n \times c}\}_{v=1}^m$ where Ψ_c denotes the shared classifier and $\hat{\mathbf{P}}_{i,j}^{(v)}$ means the confidence

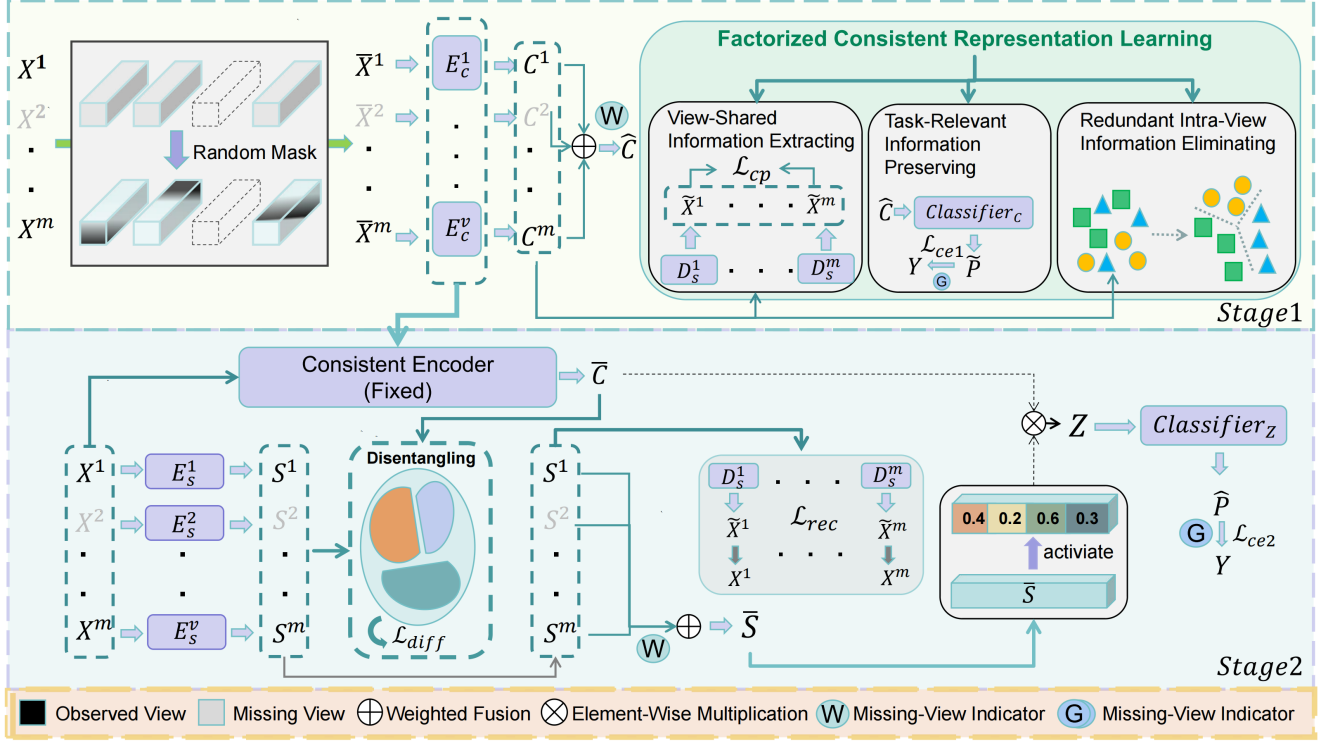


Figure 1. An overview of our MVFD framework. In the first age, we first randomly masked fragments of input features. Then we factorize consistent representation learning into three sub-goals for collaborative optimization. In the second stage, we freeze consistent encoders trained in the first stage and leverage learned consistent representation and our graph disentangling loss to guide the disentanglement process of view-specific information.




Figure 2. We define a complete multi-view data contain view-specific information, view-shared information, and redundant information between. The goal of our graph disentangling loss is fully eliminate redundancy between them and obtain disentangled view-specific information.

of sample i belongs to category j in the v -th view. In real-world situations, different views of the same instance may have different predictive confidence for the same category due to interference with view-specific information. To eliminate view-specific information and obtain pure consistent representation, we need to maintain semantic consistency between different views of the same instance. In other words, predictions from different views of the same sample should be consistent. To this end, we regard predictions as semantic

features and propose semantic contrastive loss as follows:

$$\ell_{sc} = \frac{1}{2} \sum_{v=1}^m \sum_{u \neq v}^m \ell_{sc}^{(vu)} = \frac{1}{2} \sum_{v=1}^m \sum_{u \neq v}^m \left(\sum_{i=1}^n \log \frac{\exp(\cos(\bar{\mathbf{P}}_{i,:}^{(v)}, \bar{\mathbf{P}}_{i,:}^{(u)}) / \tau)}{\exp(\cos(\bar{\mathbf{P}}_{i,:}^{(v)}, \bar{\mathbf{P}}_{i,:}^{(u)}) / \tau) + \mathcal{S}_{neg}} \right) \quad (4)$$

where $\mathcal{S}_{neg} = \sum_{r=u,v} \sum_{j=1, j \neq i}^n \exp(\cos(\bar{\mathbf{P}}_{i,:}^{(v)}, \bar{\mathbf{P}}_{j,:}^{(u)}) / \tau)$ and function $\cos(a, b)$ means the cosine similarity between embedding feature a and b . In Eq.(4), for any anchor instance $\bar{\mathbf{P}}_{i,:}^{(v)}$, we have $(N-1)$ positive instances: $\bar{\mathbf{P}}_{i,:}^{(u)} |_{u \neq v}$ and $N(M-1)$ negative instances: $\bar{\mathbf{P}}_{j,:}^{(v)} |_{j \neq i}$. Through minimizing Eq.(4), we can learn semantic consistency and eliminate view-specific information within consistent representation.

Task-relevant information preserving: Notably, the minimal sufficient representation obtained in contrastive learning is insufficient for downstream tasks since it eliminates task-relevant non-shared information [29]. To cope with this issue, inspired by [26], we can obtain task-relevant consi-

tent representation by minimizing as follows:

$$\min_{\hat{\mathbf{C}}} -I(\hat{\mathbf{C}}; \mathbf{Y}) \quad (5)$$

where $I(\hat{\mathbf{C}}; \mathbf{Y})$ means the mutual information between consistent representation $\hat{\mathbf{C}}$ and label \mathbf{Y} and minimizing Eq.(5) can be regarded as minimizing the common cross-entropy loss [8]. To achieve this goal, we take Ψ_c as classifier and put $\hat{\mathbf{C}}$ in it to obtain predictions $\tilde{\mathbf{P}} = \Psi_c(\hat{\mathbf{C}})$ where $\tilde{\mathbf{P}} \in \mathbb{R}^{n \times c}$. Then we minimize cross-entropy loss as follows:

$$\begin{aligned} \ell_{ce1} = & \frac{1}{\sum_{i,j} \mathbf{G}_{i,j}} \sum_{i=1}^n \sum_{j=1}^c \left((\mathbf{Y}_{i,j} \log(\tilde{\mathbf{P}}_{i,j}) \right. \\ & \left. + (1 - \mathbf{Y}_{i,j}) \log(1 - \tilde{\mathbf{P}}_{i,j}) \right) \mathbf{G}_{i,j} \end{aligned} \quad (6)$$

where \mathbf{G} denotes the missing-label indicator. Overall, all loss functions in the second stage becomes:

$$\ell_{stage1} = \ell_{ce1} + \alpha \ell_{cp} + \beta \ell_{sc} \quad (7)$$

3.3. View-Specific Representation Disentangling

View-specific information encodes unique semantic details within a view—such as angles, distances, and positional data. A complete view combines both view-specific and view-consistent information through complex higher-order interactions [6]. This section discusses how to disentangle these two types to obtain pure view-specific information.

An intuitive approach is to use the view consistency factor we have learned to disentangle the view-specific factor, thereby removing redundant information from the latter. Specifically, we first introduce an autoencoder to extract coarse view-specific features from the raw multi-view data: $\{E_s^{(v)} : \mathbf{X}^{(v)} \in \mathbb{R}^{n \times d_v} \rightarrow \mathbf{S}^{(v)} \in \mathbb{R}^{n \times d_e}\}_{v=1}^m$, where $E_s^{(v)}$ from $E_s^{(1)}, \dots, E_s^{(m)}$ is a set of view-specific encoders constructed by MLPs, and $\mathbf{S}^{(v)}$ represents the high-level view-specific feature of the v -th view. We then utilize a set of view-specific decoders to reconstruct the raw multi-view data: $\{D_s^{(v)} : \mathbf{S}^{(v)} \in \mathbb{R}^{n \times d_e} \rightarrow \hat{\mathbf{X}}^{(v)} \in \mathbb{R}^{n \times d_v}\}_{v=1}^m$. The reconstruction loss is calculated as follows:

$$\ell_{rec} = \frac{1}{m} \sum_{v=1}^m \ell_{rec}^{(v)} = \frac{1}{m} \sum_{v=1}^m \left(\frac{1}{d_v} \sum_{i=1}^n \left\| \hat{\mathbf{X}}_i^{(v)} - \mathbf{X}_i^{(v)} \right\|_2^2 W_{i,v} \right) \quad (8)$$

where W is the missing-view indicator, introduced to mitigate the negative effects of missing views in calculation.

Following this, the next goal is how to leverage learned consistent representation to distill pure view-specific information from coarse view-specific feature. To this end, we firstly freeze the consistent encoders trained in

the first stage. We extract shared multi-view features: $\{E_c^{(v)} : \mathbf{X}^{(v)} \in \mathbb{R}^{n \times d_v} \rightarrow \mathbf{C}^{(v)} \in \mathbb{R}^{n \times d_e}\}_{v=1}^m$. Then we can obtain a comprehensive consistent representation as follow:

$$\bar{\mathbf{C}}_{i,:} = \sum_{v=1}^m \frac{\mathbf{C}^{(v)} \mathbf{W}_{i,v}}{\sum_v \mathbf{W}_{i,v}} \quad (9)$$

Then we propose our graph disentangling loss to eliminate redundancy between different types of information as follow:

$$\begin{aligned} \ell_{gd} = & \sum_{i=1}^n \left(\frac{1}{\sum_v W_{i,v}} \sum_{v=1}^m \mathcal{S}(\bar{\mathbf{C}}_i, \mathbf{S}_i^{(v)}) W_{i,v} \right. \\ & \left. + \frac{1}{\sum_{u,v} W_{i,u} W_{i,v}} \sum_{v=1}^m \sum_{u \neq v}^m \mathcal{S}(\mathbf{S}_i^{(v)}, \mathbf{S}_i^{(u)}) W_{i,v} W_{i,u} \right) \end{aligned} \quad (10)$$

where \mathcal{S} is the cosine similarity function and $\mathcal{S}(a, b) = \frac{a^T b}{\|a\|_2 \cdot \|b\|_2}$. Equation (10) comprises two terms: the first minimizes the similarity between consistent information and view-specific information, and the second minimizes the similarity between view-specific information from different views. Specifically, $W_{i,v}$ in the first term ignores missing views, and $W_{i,v} W_{i,u}$ in the second term identifies instance pairs present in both views. By minimizing these similarities, we indirectly obtain disentangled representations while maintaining clear distinctions between view-specific features and shared features. This allows us to naturally perform weighted cross-view fusion to obtain a unique view-specific representation as follows:

$$\bar{\mathbf{S}}_{i,:} = \sum_{v=1}^m \frac{\mathbf{S}_i^{(v)} \mathbf{W}_{i,v}}{\sum_v \mathbf{W}_{i,v}} \quad (11)$$

3.4. Multi-Label Classification

Up to now, we have obtained the consistent representation $\bar{\mathbf{C}}$ and the disentangled view-specific representation $\bar{\mathbf{S}}$. The next challenge is how to combine them to derive the final discriminative feature for classification. Unlike previous works that rely on summation or concatenation operations, we are inspired by [19] and [12], which have demonstrated the effectiveness of the star operation (element-wise multiplication) in capturing subtle differences and complex patterns between data. Thus, we leverage the star operation to enhance the shared information via the view-specific information as follows:

$$\mathbf{Z} = \text{Sigmoid}(\bar{\mathbf{S}}) \odot \bar{\mathbf{C}} \quad (12)$$

where $\text{Sigmoid}(\cdot)$ is a activate function and $\mathbf{Z} \in \mathbb{R}^{n \times d_e}$ is final discriminative feature. Finally, we can obtain the final prediction $\hat{\mathbf{P}} = \Psi_Z(\mathbf{Z})$ where $\hat{\mathbf{P}} \in \mathbb{R}^{n \times c}$ and calculate

multi-label classification loss as follows:

$$\begin{aligned} \ell_{ce2} = & \frac{1}{\sum_{i,j} G_{i,j}} \sum_{i=1}^n \sum_{j=1}^c \left(\mathbf{Y}_{i,j} \log \left(\hat{\mathbf{P}}_{i,j} \right) \right. \\ & \left. + (1 - \mathbf{Y}_{i,j}) \log \left(1 - \hat{\mathbf{P}}_{i,j} \right) \right) G_{i,j} \end{aligned} \quad (13)$$

where G denotes missing-label indicator which introduced to filter invalid missing labels during calculate process.

Overall, all loss functions in the second stage becomes:

$$\ell_{stage2} = \ell_{ce2} + \gamma \ell_{rec} + \lambda \ell_{gd} \quad (14)$$

and the training process is described in Algorithm 1 in supplementary materials

4. Experiments

4.1. Experimental Settings

Datasets and evaluation metrics: Consistent with existing iMvMLC methods [13] [9] [16], we adopt five commonly datasets, i.e., Corel5k [1], Pascal07 [2], ESPGame [27], IAPRTC12 [3], and MIRFLICKR [5]. Additionally, six distinct types of features are extracted as six views from these datasets, namely GIST, HSV, HUE, LAB, RGB, and SIFT. Following previous works [13], we select six metrics to evaluate all comparison methods, i.e., Average Precision (AP), Ranking Loss (RL), the adapted area under the curve(AUC), OneError (OE), and Coverage (Cov). Specifically, for all metrics, higher values indicate better performance.

Incomplete multi-view multi-label data processing: To simulate a real-scenario of missing data, similar to existing works [23] [13], we construct incomplete multi-view multi-label datasets as follows: for each view, we randomly disable 50% instances and replace them with value '0' and ensuring that each sample has at least one available view. For each category, we randomly eliminated 50% of the positive and negative tags and replaced them with '0'. Finally, we randomly select 50% of samples with missing views and missing labels as the training set.

Comparison methods: In order to evaluate the performance of our method, we compare ten top methods, i.e., GLOCAL [39], CDMM [37], DM2L [20], LVSL [36], iMVWL [23], NAIM3L [9], DICNet [13], LMVCAT [14], MTD [12] and AIMNet [11] with our MVFD on five incomplete multi-view multi-label datasets. Specifically, not all methods for the comparison are suitable for the iMvMLC task, such as GLOCAL, CDMM, DM2L and LVSL, where GLOCAL and DM2L can deal with incomplete labels and CDMM and LVSL are the MvMLC methods unable to handle any data missing problem. To implement the above four methods on the iMvMLC tasks, we used mean imputation for missing views and regarded the unknown label as the negative tag for methods that are unable to handle partial multi-label data.

4.2. Experimental Results and Analysis

To prove the effectiveness of our method, we compare our MVFD with ten top methods on five datasets with a 50% missing rate in both views and labels, where Table 1 shows the experimental results of the six evaluation metrics. Based on the results in Table 1, we can have the following observations:

- On all five datasets, our MVFD outperforms comparison methods on all metrics, which fully verifies the effectiveness of our method.
- From the results shown in Table 1, we can observe that DNN-based methods, such as DICNet, LMVCAT, AIMNet, MTD and our MVFD, exhibit better performance than other methods, which shows the potential of deep neural network in iMvMLC task.
- The methods such as DICNet and our MVFD that consider the incompleteness problem in both views and labels, have better performance compared with other methods that either ignore the incompleteness problem in both view and label or only consider single missing problem

Additionally, to further confirm that our model has good adaptability to complete multi-view data, we provide the results of nine methods with full views and labels on Corel5k, Pascal07 and EPSGame dataset in Figure. 3, which can be observed that our MVFD still achieve competitive performance than other methods, including those designed for the ideal complete case, demonstrating the generalization ability of our model.

Moreover, to provide a more comprehensive demonstration of our model in handling missing values, we present the experimental results with different missing rates on view and label on Corel5K dataset. As shown in Fig .4a and Fig. 4b, proposed MVFD can achieve best AP score with all missing settings on View or Label.

4.3. Hyperparameters Study

The overall loss of our model mainly contains 4 hyperparameters, i.e., α , β , γ and λ . Specially, α and β are the weight coefficients of ℓ_{cp} and ℓ_{sc} in the first stage, γ and λ are the weight coefficients of ℓ_{rec} and ℓ_{gd} in the second stage. To study the impact of different hyperparameters on the performance of our model, we iteratively set different values and report the corresponding AP values on Corel5k and Pascal07 dataset, with 50% missing view and labels, with 70% training samples. Fig. 5a and Fig. 5b show the AP value versus hyperparameters α and β . Fig .5c and Fig .5d show the AP values versus hyperparameters γ and λ . As shown in Fig 5a and Fig 5b, the optimal ranges for α and β on Corel5k are [0.5,1] and [0.05,0.5], and on Pascal07, they are both

Table 1. Experimental results of eleven methods on the five datasets with 50% missing-view rate 50% missing-label rate, and 70% training samples (the bottom right digit is the standard deviation).

Data	Metric	GLOCAL	CDMM	DM2L	LVSL	iMVWL	NAIM3L	DICNet	LMVCAT	AIMNet	MTD	MVFD
Core5k	AP	0.285 _{0.004}	0.354 _{0.004}	0.262 _{0.005}	0.342 _{0.004}	0.283 _{0.008}	0.309 _{0.004}	0.381 _{0.004}	0.382 _{0.004}	0.400 _{0.010}	0.415 _{0.008}	0.441 _{0.003}
	1-HL	0.987 _{0.000}	0.987 _{0.000}	0.979 _{0.001}	0.978 _{0.000}	0.979 _{0.001}	0.988 _{0.000}	0.988 _{0.000}	0.986 _{0.000}	0.988 _{0.000}	0.980 _{0.000}	0.988 _{0.000}
	1-RL	0.840 _{0.003}	0.854 _{0.003}	0.843 _{0.002}	0.850 _{0.005}	0.878 _{0.002}	0.882 _{0.004}	0.882 _{0.003}	0.880 _{0.002}	0.902 _{0.004}	0.893 _{0.004}	0.903 _{0.003}
	AUC	0.843 _{0.003}	0.886 _{0.003}	0.845 _{0.002}	0.868 _{0.005}	0.881 _{0.002}	0.881 _{0.002}	0.884 _{0.004}	0.883 _{0.002}	0.905 _{0.003}	0.896 _{0.004}	0.908 _{0.002}
	1-OE	0.327 _{0.010}	0.410 _{0.007}	0.295 _{0.014}	0.391 _{0.009}	0.311 _{0.015}	0.350 _{0.009}	0.468 _{0.007}	0.453 _{0.006}	0.475 _{0.018}	0.491 _{0.012}	0.523 _{0.004}
	1-Cov	0.648 _{0.006}	0.723 _{0.007}	0.647 _{0.005}	0.718 _{0.006}	0.702 _{0.008}	0.725 _{0.005}	0.727 _{0.017}	0.727 _{0.006}	0.771 _{0.005}	0.749 _{0.009}	0.770 _{0.004}
Pascal07	AP	0.496 _{0.004}	0.508 _{0.005}	0.471 _{0.008}	0.504 _{0.005}	0.437 _{0.005}	0.488 _{0.003}	0.505 _{0.012}	0.519 _{0.006}	0.548 _{0.008}	0.551 _{0.004}	0.564 _{0.005}
	1-HL	0.927 _{0.000}	0.931 _{0.001}	0.928 _{0.001}	0.930 _{0.000}	0.882 _{0.004}	0.928 _{0.001}	0.929 _{0.001}	0.924 _{0.003}	0.931 _{0.001}	0.934 _{0.001}	0.935 _{0.002}
	1-RL	0.761 _{0.004}	0.812 _{0.004}	0.761 _{0.005}	0.806 _{0.003}	0.736 _{0.015}	0.783 _{0.001}	0.783 _{0.008}	0.811 _{0.004}	0.831 _{0.004}	0.832 _{0.004}	0.837 _{0.004}
	AUC	0.786 _{0.003}	0.838 _{0.003}	0.779 _{0.004}	0.832 _{0.002}	0.767 _{0.015}	0.811 _{0.001}	0.809 _{0.006}	0.834 _{0.004}	0.851 _{0.004}	0.851 _{0.003}	0.860 _{0.005}
	1-OE	0.434 _{0.005}	0.419 _{0.008}	0.410 _{0.008}	0.419 _{0.006}	0.362 _{0.023}	0.421 _{0.006}	0.427 _{0.015}	0.421 _{0.006}	0.461 _{0.013}	0.459 _{0.007}	0.474 _{0.020}
	1-Cov	0.703 _{0.004}	0.759 _{0.003}	0.692 _{0.004}	0.751 _{0.003}	0.677 _{0.015}	0.727 _{0.002}	0.731 _{0.006}	0.763 _{0.005}	0.783 _{0.004}	0.784 _{0.003}	0.790 _{0.005}
ESPGame	AP	0.221 _{0.002}	0.289 _{0.003}	0.212 _{0.002}	0.285 _{0.003}	0.244 _{0.005}	0.246 _{0.002}	0.297 _{0.002}	0.294 _{0.004}	0.305 _{0.004}	0.306 _{0.003}	0.326 _{0.003}
	1-HL	0.982 _{0.000}	0.983 _{0.000}	0.982 _{0.000}	0.983 _{0.000}	0.972 _{0.002}	0.983 _{0.000}	0.983 _{0.000}	0.982 _{0.000}	0.983 _{0.001}	0.983 _{0.000}	0.983 _{0.000}
	1-RL	0.780 _{0.004}	0.832 _{0.001}	0.781 _{0.001}	0.822 _{0.008}	0.808 _{0.002}	0.818 _{0.002}	0.828 _{0.001}	0.828 _{0.002}	0.846 _{0.002}	0.831 _{0.003}	0.848 _{0.001}
	AUC	0.784 _{0.004}	0.836 _{0.001}	0.785 _{0.001}	0.833 _{0.002}	0.813 _{0.002}	0.824 _{0.002}	0.836 _{0.001}	0.833 _{0.002}	0.850 _{0.002}	0.842 _{0.002}	0.859 _{0.002}
	1-OE	0.317 _{0.010}	0.389 _{0.004}	0.313 _{0.009}	0.389 _{0.004}	0.343 _{0.013}	0.339 _{0.003}	0.439 _{0.005}	0.434 _{0.009}	0.442 _{0.011}	0.447 _{0.009}	0.480 _{0.005}
	1-Cov	0.490 _{0.006}	0.574 _{0.004}	0.488 _{0.005}	0.567 _{0.005}	0.548 _{0.004}	0.571 _{0.003}	0.593 _{0.003}	0.590 _{0.004}	0.624 _{0.005}	0.602 _{0.004}	0.633 _{0.005}
IAPRTC12	AP	0.256 _{0.002}	0.305 _{0.004}	0.234 _{0.003}	0.304 _{0.004}	0.237 _{0.003}	0.261 _{0.003}	0.323 _{0.003}	0.317 _{0.003}	0.329 _{0.005}	0.332 _{0.003}	0.360 _{0.003}
	1-HL	0.980 _{0.000}	0.981 _{0.000}	0.980 _{0.000}	0.981 _{0.000}	0.969 _{0.002}	0.980 _{0.000}	0.981 _{0.000}	0.980 _{0.000}	0.981 _{0.001}	0.981 _{0.000}	0.981 _{0.000}
	1-RL	0.832 _{0.004}	0.862 _{0.002}	0.823 _{0.003}	0.861 _{0.002}	0.833 _{0.002}	0.844 _{0.003}	0.873 _{0.001}	0.870 _{0.001}	0.883 _{0.003}	0.875 _{0.002}	0.888 _{0.005}
	AUC	0.830 _{0.001}	0.864 _{0.002}	0.825 _{0.003}	0.863 _{0.003}	0.835 _{0.001}	0.850 _{0.001}	0.874 _{0.002}	0.872 _{0.001}	0.885 _{0.003}	0.876 _{0.001}	0.896 _{0.001}
	1-OE	0.430 _{0.008}	0.429 _{0.005}	0.352 _{0.008}	0.429 _{0.005}	0.314 _{0.009}	0.386 _{0.003}	0.468 _{0.007}	0.443 _{0.005}	0.459 _{0.080}	0.467 _{0.005}	0.507 _{0.007}
	1-Cov	0.534 _{0.003}	0.597 _{0.004}	0.529 _{0.004}	0.597 _{0.004}	0.564 _{0.005}	0.592 _{0.004}	0.649 _{0.004}	0.648 _{0.003}	0.673 _{0.001}	0.649 _{0.004}	0.688 _{0.007}
MIRFlickr	AP	0.537 _{0.005}	0.570 _{0.006}	0.514 _{0.004}	0.553 _{0.004}	0.490 _{0.010}	0.510 _{0.005}	0.589 _{0.005}	0.594 _{0.005}	0.602 _{0.009}	0.607 _{0.004}	0.615 _{0.001}
	1-HL	0.874 _{0.001}	0.888 _{0.001}	0.882 _{0.001}	0.885 _{0.001}	0.839 _{0.002}	0.882 _{0.001}	0.888 _{0.001}	0.882 _{0.002}	0.890 _{0.001}	0.891 _{0.001}	0.894 _{0.001}
	1-RL	0.832 _{0.001}	0.844 _{0.002}	0.844 _{0.002}	0.856 _{0.004}	0.803 _{0.004}	0.844 _{0.002}	0.830 _{0.007}	0.863 _{0.004}	0.883 _{0.003}	0.875 _{0.002}	0.890 _{0.004}
	AUC	0.817 _{0.002}	0.844 _{0.001}	0.830 _{0.001}	0.844 _{0.001}	0.787 _{0.003}	0.837 _{0.001}	0.849 _{0.002}	0.853 _{0.003}	0.861 _{0.001}	0.862 _{0.001}	0.873 _{0.003}
	1-OE	0.511 _{0.012}	0.579 _{0.002}	0.517 _{0.008}	0.631 _{0.002}	0.490 _{0.010}	0.510 _{0.005}	0.652 _{0.007}	0.642 _{0.008}	0.651 _{0.006}	0.655 _{0.004}	0.672 _{0.005}
	1-Cov	0.605 _{0.004}	0.649 _{0.003}	0.597 _{0.003}	0.636 _{0.004}	0.572 _{0.013}	0.631 _{0.002}	0.652 _{0.007}	0.667 _{0.003}	0.671 _{0.004}	0.676 _{0.004}	0.690 _{0.002}




Figure 3. Experimental results of eleven methods on three datasets without any missing views or labels. The worst results are indicated at the center of the radar chart while the best results are represented by the vertices on the six metrics.

[0.05,0.1]. From Fig. 5c and Fig. 5d, the optimal ranges for γ and λ on Core5k are [0.1,0.5] and [0.005,0.01], and on Pascal07, they are [0.05,1] and [0.05,0.5], respectively.

4.4. Analysis of Disentanglement

To demonstrate the effectiveness of our graph disentangling loss, inspired by [12], we plot a random sample’s channel similarity heat maps across view-shared feature and view-specific feature in different training epochs in stage

2 on the Core5k dataset with half of missing views and labels. For any sample i , we obtain \mathbf{O}_i by concatenating its view-shared feature $\left\{ \mathbf{C}_i^{(v)} \right\}_{v=1}^m$ with its view-specific feature $\left\{ \mathbf{S}_i^{(v)} \right\}_{v=1}^m$, then we can calculate the i -th sample’s heat map $\mathbf{H}_i = \mathbf{O}_i \mathbf{O}_i^T$ where \mathbf{H}_i shows the cosine similarity of sample i ’s different features. From Fig. 6(a), we can observe that in the initial state of stage 2, the similarities for the view-shared features are high, which demonstrate the effectiveness

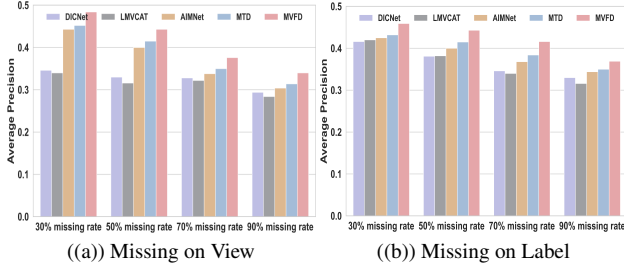


Figure 4. Performance about the comparison with SOTA methods with different missing rates on View or Label




Figure 5. The AP values for hyper-parameters α and β on the Corel5k (Fig5. a) and Pascal07 (Fig5. b) datasets; AP values for hyper-parameters α and β on the Corel5k (Fig5. c) and Pascal07 (Fig5. d) datasets are presented. Both datasets contain 50% available views and labels, with a 70% training sample rate.

of our factorized consistent representation learning strategy in stage 1. From Fig. 6(b)(c)(d), it is evident that as training progresses, The similarity between view-consistent features and view-specific features, as well as between different view specific features, is gradually decrease, which indicates that our graph disentangling loss is functioning as expected.

Table 2. The ablation experiment result on Corel5k dataset and Pascal07 dataset.

Backbone	L_{cp}	L_{sc}	L_{rec}	L_{gd}	Corel5k		Pascal07	
					AP	AUC	AP	AUC
✓					0.398	0.890	0.530	0.820
✓	✓				0.416	0.899	0.545	0.845
✓		✓			0.415	0.897	0.543	0.844
✓			✓		0.403	0.892	0.534	0.833
✓				✓	0.406	0.898	0.538	0.840
✓	✓				0.430	0.904	0.565	0.851
✓		✓			0.429	0.905	0.563	0.854
✓	✓	✓			0.441	0.908	0.564	0.860
one stage					0.426	0.900	0.547	0.845

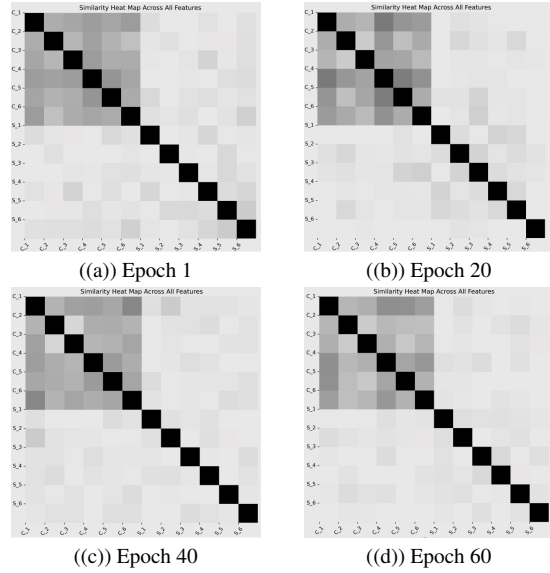


Figure 6. A random sample’s feature similarity heat maps across different informations on Corel5k dataset.

4.5. Ablation Study

To validate each module’s effectiveness, we conduct ablation experiments on Corel5k and Pascal07 datasets with a 50% missing rate of views and labels. Loss functions ℓ_{cp} , ℓ_{sc} , ℓ_{rec} , and ℓ_{gd} are sequentially removed, while the backbone retains classification losses ℓ_{ce1} and ℓ_{ce2} in two stages. To compare the two-stage learning paradigm with an end-to-end framework, we consolidate the two-stage model into an end-to-end one by moving first-stage losses to the second stage and training all modules simultaneously. Results in Table 2 show: (i): All designs introduced into our method are crucial and result in an enhancement in performance metrics. (ii): When we transfer the two-stage model to the end-to-end model, there is a performance degradation, which confirms that the two-stage learning paradigm may have an advantage in extracting consistent and view-specific information from multi-view data compared with the end-to-end framework.

5. Conclusion

In this paper, we propose a novel framework, namely MVFD for incomplete multi-view multi-label classification (iMvMLC) tasks. Specifically, we novelty factorize consistency learning into three sub-objects and design corresponding collaborative consistency learning modules based on masked cross-view prediction (MCP) strategy and information theory. Furthermore, we decompose multi-view representation into view-consist representation and view-specific representation, and proposed a graph disentangling loss to fully eliminate redundant information between them. Finally, we conduct extensive comparison and ablation experiments to demonstrate the effectiveness of our method.

References

- [1] P. Duygulu, K. Barnard, J. F. G. de Freitas, and D. A. Forsyth. Object recognition as machine translation-i : Learning a lexicon for a fixed image vocabulary. 2002.
- [2] Mark Everingham, S. M. Ali Eslami, Luc Van Gool, Christopher K. I. Williams, John M. Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision*, 111:98 – 136, 2014.
- [3] Michael Grubinger, Paul D. Clough, Henning Müller, and Thomas Deselaers. The iapr tc-12 benchmark: A new evaluation resource for visual information systems. 2006.
- [4] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. Masked autoencoders are scalable vision learners. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15979–15988, 2021.
- [5] Mark J. Huiskes and Michael S. Lew. The mir flickr retrieval evaluation. In *Multimedia Information Retrieval*, 2008.
- [6] Guanzhou Ke, Bo Wang, Xiaoli Wang, and Shengfeng He. Rethinking multi-view representation learning via distilled disentangling. *ArXiv*, abs/2403.10897, 2024.
- [7] Takumi Kobayashi. for two-way multi-label loss. 2023.
- [8] Changhee Lee and Mihaela van der Schaar. A variational information bottleneck approach to multi-omics data integration. In Arindam Banerjee and Kenji Fukumizu, editors, *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 1513–1521. PMLR, 13–15 Apr 2021.
- [9] Xiang Li and Songcan Chen. A concise yet effective model for non-aligned incomplete multi-view and missing multi-label learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10):5918–5932, 2022.
- [10] Yawen Ling, Jianpeng Chen, Yazhou Ren, Xiaorong Pu, Jie Xu, Xiao lan Zhu, and Lifang He. Dual label-guided graph refinement for multi-view graph clustering. In *AAAI Conference on Artificial Intelligence*, 2023.
- [11] Chengliang Liu, Jinlong Jia, Jie Wen, Yabo Liu, Xiaoling Luo, Chao Huang, and Yong Xu. Attention-induced embedding imputation for incomplete multi-view partial multi-label classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 13864–13872, 2024.
- [12] Chengliang Liu, Jie Wen, Yabo Liu, Chao Huang, Zhihao Wu, Xiaoling Luo, and Yong Xu. Masked two-channel decoupling framework for incomplete multi-view weak multi-label learning. In *Advances in Neural Information Processing Systems*, volume 36, 2024.
- [13] Chengliang Liu, Jie Wen, Xiaoling Luo, Chao Huang, Zhihao Wu, and Yong Xu. Dicnet: Deep instance-level contrastive network for double incomplete multi-view multi-label classification. *ArXiv*, abs/2303.08358, 2023.
- [14] Chengliang Liu, Jie Wen, Xiaoling Luo, and Yong Xu. Incomplete multi-view multi-label learning via label-guided masked view- and category-aware transformers. *ArXiv*, abs/2303.07180, 2023.
- [15] Meng Liu, Yong Luo, Dacheng Tao, Chao Xu, and Yonggang Wen. Low-rank multi-view learning in matrix completion for multi-label image classification. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI’15, page 2778–2784. AAAI Press, 2015.
- [16] Xinwang Liu, Xinzhong Zhu, Miaomiao Li, Lei Wang, Chang Tang, Jianping Yin, Dinggang Shen, Huaimin Wang, and Wen Gao. Late fusion incomplete multi-view clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41:2410–2423, 2019.
- [17] Jiang Long, Qi Zhang, Xiaohuan Lu, Jie Wen, Lian Zhao, and Wulin Xie. Multi-scale locality preserving projection for partial multi-view incomplete multi-label learning. *Neural networks : the official journal of the International Neural Network Society*, 180:106748, 2024.
- [18] Xiaohuan Lu, Lian Zhao, Wai Keung Wong, Jie Wen, Jiang Long, and Wulin Xie. Task-augmented cross-view imputation network for partial multi-view incomplete multi-label classification. *ArXiv*, abs/2409.07931, 2024.
- [19] Xu Ma, Xiyang Dai, Yue Bai, Yizhou Wang, and Yun Fu. Rewrite the stars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.
- [20] Zhongchen Ma and Songcan Chen. Expand globally, shrink locally: Discriminant multi-label learning with missing labels. *Pattern Recognit.*, 111:107675, 2020.
- [21] Bingyan Nie, Wulin Xie, Jiang Long, and Xiaohuan Lu. Incomplete multi-view multi-label classification via a dual-level contrastive learning framework, 2024.
- [22] Shiliang Sun and Daoming Zong. Lcbm: A multi-view probabilistic model for multi-label classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(8):2682–2696, 2021.
- [23] Qiaoyu Tan, Guoxian Yu, Carlotta Domeniconi, Jun Wang, and Zi Jian Zhang. Incomplete multi-view weak-label learning. In *International Joint Conference on Artificial Intelligence*, 2018.
- [24] Yuze Tan, Yixi Liu, Shudong Huang, Wentao Feng, and Jiancheng Lv. Sample-level multi-view graph clustering. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 23966–23975, 2023.
- [25] Jidapa Thadajarassiri, Thomas Hartvigsen, Walter Gerych, Xiangnan Kong, and Elke Rundensteiner. Knowledge amalgamation for multi-label classification via label dependency transfer. In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI’23/IAAI’23/EAAI’23. AAAI Press, 2023.
- [26] Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. *ArXiv*, physics/0004057, 2000.
- [27] Luis von Ahn and Laura A. Dabbish. Labeling images with a computer game. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2004.
- [28] Xinhang Wan, Xinwang Liu, Jiyuan Liu, Siwei Wang, Yi Wen, Weixuan Liang, En Zhu, Zhe Liu, and Lu Zhou. Auto-weighted multi-view clustering for large-scale data. In *AAAI Conference on Artificial Intelligence*, 2023.
- [29] Haoqing Wang, Xun Guo, Zhiwei Deng, and Yan Lu. Rethinking minimal sufficient representation in contrastive learning.

2022 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16020–16029, 2022.

- [30] Xuan Wu, Qing-Guo Chen, Yao Hu, Dengbao Wang, Xiaodong Chang, Xiaobo Wang, and Min-Ling Zhang. Multi-view multi-label learning with view-specific information extraction. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence, IJCAI'19*, page 3884–3890. AAAI Press, 2019.
- [31] Wulin Xie, Xiaohuan Lu, Yadong Liu, Jiang Long, Bob Zhang, Shuping Zhao, and Jie Wen. Uncertainty-aware pseudo-labeling and dual graph driven network for incomplete multi-view multi-label classification. In *ACM Multimedia*, 2024.
- [32] Cai Xu, Jiajun Si, Ziyu Guan, Wei Zhao, Yue Wu, and Xiyue Gao. Reliable conflictive multi-view learning. In *AAAI Conference on Artificial Intelligence*, 2024.
- [33] Jie Xu, Huayi Tang, Yazhou Ren, Liang Peng, Xiao lan Zhu, and Lifang He. Multi-level feature learning for contrastive multi-view clustering. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16030–16039, 2021.
- [34] Changqing Zhang, Yajie Cui, Zongbo Han, Joey Tianyi Zhou, Huazhu Fu, and Qinghua Hu. Deep partial multi-view learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(5):2402–2415, 2022.
- [35] Changqing Zhang, Ziwei Yu, Qinghua Hu, Pengfei Zhu, Xinwang Liu, and Xiaobo Wang. Latent semantic aware multi-view multi-label classification. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence, AAAI'18/IAAI'18/EAAI'18*. AAAI Press, 2018.
- [36] Dawei Zhao, Qingwei Gao, Yixiang Lu, and Dong Sun. Non-aligned multi-view multi-label classification via learning view-specific labels. *IEEE Transactions on Multimedia*, 25:7235–7247, 2023.
- [37] Dawei Zhao, Qingwei Gao, Yixiang Lu, Dong Sun, and Yusheng Cheng. Consistency and diversity neural network multi-view multi-label learning. *Knowl. Based Syst.*, 218:106841, 2021.
- [38] Changming Zhu, Duoqian Miao, Zhe Wang, Rigui Zhou, Lai Wei, and Xiafen Zhang. Global and local multi-view multi-label learning. *Neurocomput.*, 371(C):67–77, jan 2020.
- [39] Yue Zhu, James T. Kwok, and Zhi-Hua Zhou. Multi-label learning with global and local label correlation. *IEEE Transactions on Knowledge and Data Engineering*, 30(6):1081–1094, 2018.