

RAC-DMVC: Reliability-Aware Contrastive Deep Multi-View Clustering Under Multi-Source Noise

Shihao Dong¹, Yue Liu², Xiaotong Zhou¹, Yuhui Zheng^{3*}, Huiying Xu⁴, Xinzhong Zhu⁴

¹School of Computer Science, Nanjing University of Information Science and Technology, Nanjing, China.

²School of Computing, National University of Singapore, Singapore.

³Key Laboratory of Tibetan Information Processing, Ministry of Education, Qinghai Normal University. Xining, China.

⁴School of Computer Science and Technology, Zhejiang Normal University. Jinhua, China.

{dongshihao, xiaotong_zhou}@nuist.edu.cn, yliu@u.nus.edu, zhengyh@vip.126.com, {xhy, zxz}@zjnu.edu.cn

Abstract

Multi-view clustering (MVC), which aims to separate the multi-view data into distinct clusters in an unsupervised manner, is a fundamental yet challenging task. To enhance its applicability in real-world scenarios, this paper addresses a more challenging task: MVC under multi-source noises, including missing noise and observation noise. To this end, we propose a novel framework, Reliability-Aware Contrastive Deep Multi-View Clustering (RAC-DMVC), which constructs a reliability graph to guide robust representation learning under noisy environments. Specifically, to address observation noise, we introduce a cross-view reconstruction to enhance robustness at the data level, and a reliability-aware noise contrastive learning to mitigate bias in positive and negative pairs selection caused by noisy representations. To handle missing noise, we design a dual-attention imputation to capture shared information across views while preserving view-specific features. In addition, a self-supervised cluster distillation module further refines the learned representations and improves the clustering performance. Extensive experiments on five benchmark datasets demonstrate that RAC-DMVC outperforms SOTA methods on multiple evaluation metrics and maintains excellent performance under varying ratios of noise.

Code — <https://github.com/LouisDong95/RAC-DMVC>

Introduction

In recent years, advances in sensing and data acquisition technologies have facilitated the widespread use of multi-modal and multi-view data across a variety of practical applications. Multi-view data captures complementary information from heterogeneous sources (e.g., images, text, signals) and offers multiple perspectives on the same entity, thereby enhancing data expressiveness. Effectively integrating such information has become critical in domains such as industrial inspection (Wang et al. 2023; Asad et al. 2025), social network analysis (Chen, Chen, and Gan 2024; Huang et al. 2018), healthcare (Li and Zhou 2025; Holm et al. 2025), and biomedicine (Cui et al. 2025; Rao et al. 2025).

*Corresponding author

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

Consequently, clustering multi-view data to uncover its latent structure and patterns has attracted increasing attention from the research community.

As a key technique for analyzing multi-view data, MVC aims to leverage complementary information from diverse perspectives or modalities to perform unsupervised clustering. Early MVC methods were primarily rooted in traditional machine learning approaches, including subspace learning (Gao et al. 2015; Cao et al. 2015; Kang et al. 2020), non-negative matrix factorization (Liu et al. 2013; Zhao, Ding, and Fu 2017; Huang, Kang, and Xu 2020), and graph-based learning (Tang et al. 2020; Wang, Yang, and Liu 2019). These methods promote the effective fusion of multi-view information by designing consistency constraints or shared latent representations. With the advent of deep learning, deep multi-view clustering (DMVC) (Li et al. 2019; Ren et al. 2024; Zhang et al. 2025a,b) has emerged, offering superior performance under complex data distributions through powerful end-to-end feature learning. Although multi-view clustering methods have achieved remarkable progress, they still encounter significant challenges caused by noise in real-world applications.

In the acquisition and processing of multi-view data, the presence of noise is inevitable—particularly view missing, which often results from acquisition errors, sensor failures, or communication interruptions. Existing methods for handling missing views include imputation-based methods (Tang and Liu 2022; Jin et al. 2023; Li et al. 2023; Pu et al. 2024; Yuan et al. 2025; Chao, Jiang, and Chu 2024) which aim to restore the missing views using available information, and non-imputation methods (Lin et al. 2023; Xu et al. 2022; Lu et al. 2024; Feng et al. 2024), which avoid the additional noise caused by imputation through cross-view prediction. Although existing methods perform well in handling missing noise, they often overlook observation noise, such as lighting artifacts in nighttime images, background interference in audio signals, or blurred handwriting in text data. In real-world scenarios, observation noise is not only more prevalent than missing noise but also more likely to be overlooked. Noisy views often fail to provide useful information and may even degrade the quality of fused features, posing a significant challenge to multi-view clustering. Therefore, developing a unified and robust framework

to jointly handle both missing and observation noise has become a crucial step toward enhancing the reliability and applicability of MVC in practical noisy environments.

To effectively address the challenge of multi-source noise in MVC, we propose a Reliability-Aware Contrastive Deep Multi-View Clustering framework. This framework constructs a reliability graph based on inter-sample similarity to guide noise contrastive learning and dual-attention imputation to mitigate the adverse effects of observation noise and missing noise on clustering performance. Specifically, we first cross-view reconstruct clean view from noisy views, enhancing robustness at the data level. Then, a feature similarity graph is constructed to simultaneously guide the contrastive learning against observation noise and the imputation process for missing data. For observation noise, the graph serves as the weights of positive and negative pairs for reliability-aware noise contrastive to mitigate the bias in positive and negative pair selection caused by noise. For missing noise, we design a dual-attention imputation to jointly model view-specific features and cross-view shared representations to assign adaptive attention weights, thereby improving imputation accuracy. Finally, we incorporate a self-supervised cluster distillation strategy to leverage the cluster distribution information to regularize and guide the learning of view representation learning. This collaborative design significantly improves the clustering performance under noisy scenarios. Our main contributions are as follows:

- For the first time, we systematically consider the impact of missing and observation noise in the MVC task and propose a reliability-aware contrastive deep multi-view clustering method.
- By constructing a reliable graph to guide noise contrastive learning and dual-attention imputation, our method effectively mitigates the effects of both observation and missing noise.
- On five widely used multi-view datasets, the performance of our method under various noise ratios is systematically evaluated, and the results fully verify the superiority of our method compared with the SOTA methods.

Related Works

Deep Multi-view Clustering

In recent years, DMVC has been greatly developed due to the powerful feature extraction ability of deep neural networks. According to the clustering method can be divided into three categories: 1) DEC-based: These approaches are typically built upon Deep Embedded Clustering, which jointly optimizes feature learning and clustering in a unified framework. DAMC (Li et al. 2019) proposes a deep adversarial MVC network to learn the intrinsic structure embedding in multi-view data. AIMC (Xu et al. 2019) extends it to handle incomplete multi-view data. Multi-VAE (Xu et al. 2021) proposes a VAE-based MVC framework by learning disentangled visual representations. 2) Subspace clustering-based: These methods aim to project multi-view data into a common latent subspace

where self-expressiveness holds. DMSC (Abavisani and Patel 2018) inserts self-expression layers in AEs and compares the clustering performance of the fusion methods at different stages. DMVSSC (Tang et al. 2018) proposes a deep multi-view sparse subspace clustering model consisting of convolutional-AE and self-expressive module. 3) GNN-based: These methods utilize graph neural networks to model and aggregate multi-view relational information through graph structures. MAGCN (Cheng et al. 2021) proposes a multi-view attribute graph convolution network model to handle node attributes and graph reconstruction. DFMVC (Ren et al. 2024) proposes a dynamic graph fusion method to guide the learning of each view feature through the consistent features fused.

Noisy Multi-view Learning

Multi-view learning faces the challenge of multi-source noise in real-world scenarios, mainly including missing noise and non-missing noise. For missing noise, existing research is relatively mature, including: 1) Imputation methods (Tang and Liu 2022; Yang et al. 2023; Jin et al. 2023) use observable neighboring samples to achieve imputation, thereby retaining isomorphic structural information; Some works (Li et al. 2023; Pu et al. 2024; Yuan et al. 2025) perform imputation by calculating reliable prototypes to enhance the discriminability of imputation results. 2) Prediction methods (Xu et al. 2022; Lin et al. 2023; Lu et al. 2024) adopt cross-view prediction strategies to avoid additional noise introduced by the interpolation process. In addition, DMVG (Wen et al. 2024) introduces a conditional diffusion model to use available views as conditions to guide the generation of missing views. For non-missing noise, it mainly includes observation noise (Xu et al. 2024; Yang et al. 2025), pseudo-label noise (Sun et al. 2025) and misalignment noise (Yang et al. 2023; Guo et al. 2024): For observation noise, MVMAN (Xu et al. 2024) combines feature representation with soft labels to optimize robust targets to improve noise resistance; AIRMVC (Yang et al. 2025) models noise recognition as anomaly detection tasks, introduces gaussian mixture models for anomaly modeling, and weakens the impact of noise through a hybrid rectification strategy, while combining contrastive learning to obtain a more robust feature representation.

Methods

Preliminaries

Given a multi-view dataset $\mathcal{X} = \{\mathbf{X}^1, \dots, \mathbf{X}^V\}$ of N samples V views, $\mathbf{X}^v = \{\mathbf{x}_1^v, \dots, \mathbf{x}_N^v\} \in \mathbb{R}^{d_v \times N}$, where \mathbf{x}_i^v represents the i -th sample of the v -th view, and d_v represents the dimension of the v -th view data. Due to the presence of noise, the multi-view data is divided into normal data $\mathbf{X}_{\text{normal}}$ and noise data $\mathbf{X}_{\text{noisy}}$, $\mathbf{X} = \mathbf{X}_{\text{normal}} \cup \mathbf{X}_{\text{noisy}}$. Due to the existence of multi-source noise, the noise includes missing noise $\bar{\mathbf{X}}$ and observation noise $\tilde{\mathbf{X}}$, $\mathbf{X}_{\text{noisy}} = \bar{\mathbf{X}} \cup \tilde{\mathbf{X}}$. Let $\eta_m, \eta_n \in [0, 1]$ denote the ratios of missing noise and observation noise, respectively. Our objective is to divide the unlabeled data into K clusters when there is complex noise in the multi-view data, so that the samples in




Figure 1: **Overview of the our pipeline.** The processing flow begins with cross-view reconstruction for robust representation. Subsequently, a similarity-based reliability graph is constructed to guide both contrastive learning and a dual-attention imputation module. Finally, self-supervised clustering distillation is applied to refine the view-specific representations.

the same cluster are closer and the samples in different clusters are farther away. The pipeline of our method is shown in the Fig. 1.

Cross-view Reconstruction for Denoising

In multi-view learning, heterogeneous data typically contain both view-specific features and shared correlation information across views. Existing methods often employ intra-view reconstruction via autoencoders to capture view-specific representations. However, this strategy suffers from two key limitations in noisy scenarios. 1) Overfitting to local noise: Noisy observations within a view are directly reconstructed, leading to corrupted latent features. 2) Lack of semantic integration: Independent view-wise encoding overlooks the cross-view complementarity and global consistency. To address these issues, we propose a cross-view reconstruction to guide the encoded features of one view to reconstruct the data from another view:

$$\ell_{\text{crec}}(u, v) = \left\| \mathbf{X}^v - \mathcal{D}^v \left(\mathcal{E}^u(\tilde{\mathbf{X}}^u) \right) \right\|_2^2, \quad (1)$$

$$\ell_{\text{crec}} = \sum_{\substack{u, v \\ v \neq u}}^V \ell_{\text{crec}}(u, v). \quad (2)$$

Among them, \mathcal{E}^u , \mathcal{D}^v represent the u -th encoder and the v -th decoder respectively. Through the cross-view reconstruction mechanism, not only can the observation noise be effectively fitted, but also the view-specific features can be retained while incorporating information between different views.

Reliability-aware Noise Contrastive Learning

Traditional multi-view contrastive learning methods usually regard the representations of the same sample under differ-

ent views as positive pairs, and other sample pairs as negative pairs, so as to pull positive pairs closer and push negative pairs further away:

$$\ell_{\text{con}} = -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp(\mathbf{z}_i^u \top \mathbf{z}_i^v / \tau)}{\sum_{j=1}^N \exp(\mathbf{z}_i^u \top \mathbf{z}_j^v / \tau)}. \quad (3)$$

However, this hard segmentation strategy inevitably produces false negative pairs during the sample pair construction process, i.e., samples of the same type are regarded as negative pairs, which significantly reduces the quality of feature representation. To address this issue, recent methods (Lu et al. 2024; Guo et al. 2024) use similarity as weights to calculate additional positive pairs. While this approach improves the diversity of positive pairs and alleviates the issue of false negatives, the problem of false positive pairs remains unresolved. Moreover, the selection of positive and negative pairs is not reliable due to the influence of noisy representation. This instability ultimately degrades the robustness and generalization of MVC. To address these limitations, we propose a reliability-aware noise contrastive learning that introduces a reliability graph to model inter-sample relationships and guide sample pair selection. The reliability graph is constructed as follows:

$$s_{ij}^{uv} = \begin{cases} \exp\left(-\frac{\|\mathbf{z}_i^u - \mathbf{z}_j^v\|^2}{\sigma}\right), & \text{if } i \neq j \\ 1, & \text{if } i = j \end{cases} \quad (4)$$

$$\mathbf{A}^{uv} = \mathbf{S} \mathbf{D}^{-1}. \quad (5)$$

where, σ is scaling factor, s_{ij}^{uv} represents the similarity between the i -th sample of the u -th view and the j -th sample of the v -th view. \mathbf{D} is degree matrix, $d_{ii} = \sum_j s_{ij}$, a_{ij} represents the probability that sample i and sample j are judged

as positive pairs. By converting the similarity information between samples into probability weights, similarity is used as a measure of the reliable relationship between samples to guide noise contrastive learning:

$$\ell_{\text{ncon}}(u, v) = - \sum_{i=1}^N \log \frac{\exp(\sum_{j=1}^N a_{ij}^{uv} \cdot \mathbf{z}_i^u \top \mathbf{z}_j^v / \tau_c)}{\sum_{j=1}^N (1 - a_{ij}^{uv}) \cdot \exp(\mathbf{z}_i^u \top \mathbf{z}_j^v / \tau_c)}. \quad (6)$$

Where a_{ij} denote the reliability-aware indicator of a positive pair, where $(1 - a_{ij})$ naturally reflects a negative pair. τ_c denotes the temperature parameter in contrastive learning. By incorporating the reliability weights into contrastive training, the model can adaptively emphasize trustworthy sample pairs and downweight noisy or uncertain ones. This strategy not only improves the robustness of contrastive objectives under noisy conditions but also enhances the discriminative capacity of the learned representations by preserving reliable semantic structures. The final noise contrastive loss is defined as follows:

$$\ell_{\text{ncon}} = \frac{1}{N} \sum_{\substack{u, v \\ v \neq u}}^V (\ell_{\text{ncon}}(u, v) + \ell_{\text{ncon}}(u, u)). \quad (7)$$

The noise contrastive loss consists of inter-view contrastive and intra-view contrastive.

Dual-attention Imputation

To reduce the impact of missing samples on model performance, we propose a method that combines intra-view and inter-view dual attention imputation. Specifically, for missing samples in u -th view, we use the corresponding samples in v -th view as queries, calculate their similarity with the observable samples in u -th view and the observable samples in v -th view. These observable samples are then treated as keys and values in an attention module to perform weighted imputation, thereby estimating the missing representations. The inter-view and intra-view attention weights are computed as follows:

$$\mathbf{S}_{\text{inter}} = \text{Softmax}((2\mathbf{Z}_{[m^u]}^v \top \mathbf{Z}_{[-m^u]}^u - 2)/\sigma), \quad (8)$$

$$\mathbf{S}_{\text{intra}} = \text{Softmax}((2\mathbf{Z}_{[m^u]}^v \top \mathbf{Z}_{[-m^v]}^v - 2)/\sigma). \quad (9)$$

Where, $[m^u]$, $[-m^u]$ represent the missing index and observable sample index of the u -th view respectively. $\mathbf{Z}_{[m^u]}^v$ indicates that the samples corresponding to v -th view are missing in u -th view. $\mathbf{Z}_{[-m^u]}^u$, $\mathbf{Z}_{[-m^v]}^v$ represent the observable samples of views u and v respectively. $\mathbf{S}_{\text{inter}}$, $\mathbf{S}_{\text{intra}}$ represent the inter-view and intra-view attention relations respectively. The final imputation of missing samples in the u -th view is given by:

$$\mathbf{Z}_{[m^u]}^u = \alpha \cdot \mathbf{Z}_{[-m^u]}^u \mathbf{S}_{\text{inter}} + (1 - \alpha) \cdot \mathbf{Z}_{[-m^v]}^v \mathbf{S}_{\text{intra}}. \quad (10)$$

This strategy not only retains the homogeneous structural information within the view, but also effectively integrates the heterogeneous completion information across views, thereby achieving a more robust estimation of missing features.

Self-supervised Cluster Distillation

In order to more effectively utilize the shared information to guide view-specific feature learning, we designed a self-supervised cluster distillation to constrain the distribution relationship between the fusion features and the view-specific features. Specifically, the view-specific features $\mathbf{Z}^v \in \mathbb{R}^{d \times N}$ are first concatenated to form a fusion feature matrix $\mathbf{Z} \in \mathbb{R}^{D \times N}$, $D = Vd$. Then, the fused features are clustered with k -means to obtain pseudo labels and calculate K cluster centers $\mathbf{C} \in \mathbb{R}^{D \times K}$:

$$\mathbf{c}_k = \frac{1}{|\mathcal{M}_k|} \sum_{i \in \mathcal{M}_k} \mathbf{z}_i. \quad (11)$$

where, \mathcal{M}_k represents the feature set with pseudo label k . For each fused feature \mathbf{z}_i , the soft target distribution \mathbf{Q} of cluster assignment is defined as follows:

$$q_{ij} = \frac{\exp(\mathbf{z}_i \top \mathbf{c}_j / \tau_d)}{\sum_{k=1}^K \exp(\mathbf{z}_i \top \mathbf{c}_k / \tau_d)}. \quad (12)$$

Where, q_{ij} represents the probability that the i -th feature belongs to the j -th cluster. τ_d represents the temperature parameter, which is used to control the sharpness of the distribution.

For each view, we introduce a trainable clustering layer, use randomly initialized parameters $\boldsymbol{\mu}^v$ as cluster centers, and calculate the predicted distribution \mathbf{P}^v between view features and view cluster centers:

$$p_{ij}^v = \frac{\exp(\mathbf{z}_i^v \top \boldsymbol{\mu}_j^v / \tau_d)}{\sum_{k=1}^K \exp(\mathbf{z}_i^v \top \boldsymbol{\mu}_k^v / \tau_d)}. \quad (13)$$

Finally, the cluster distillation loss is defined as the sum of the Kullback-Leibler divergence between the target distribution \mathbf{Q} and the prediction distribution \mathbf{P}^v for each view:

$$\ell_{\text{dist}} = \sum_{v=1}^V \text{KL}(\mathbf{P}^v \parallel \mathbf{Q}). \quad (14)$$

To summarize, the total loss of our method consists of cross-view reconstruction loss, noise contrastive loss, and distillation loss:

$$\ell = \ell_{\text{rec}} + \ell_{\text{ncon}} + \ell_{\text{dist}}. \quad (15)$$

Experiments

Datasets and Evaluation Metrics

- **Scene-15** (Fei-Fei and Perona 2005) includes 4,485 images across 15 categories. We employ PHOG and GIST as two distinct views following (Yang et al. 2023).
- **Caltech-101** (Li et al. 2015) consists 8,677 images collected from 101 classes. We use two kinds of deep features extracted by the DECAF and VGG19 neural networks as two views following (Han et al. 2021).
- **LandUse-21** (Yang and Newsam 2010) contains 2,100 satellite imagery samples in 21 categories. We employ the PHOG and LBP features as two views following (Lin et al. 2023).

Ratio	Methods	Scene15			Caltech101			LandUse21			Reuters			NUS-WIDE		
		ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI
0%	DCCAE (Wang et al. 2015)	34.6	39.0	19.7	45.8	68.6	37.7	15.6	24.4	4.4	42.0	20.3	8.5	47.5	17.1	37.6
	DSIMVC (Tang and Liu 2022)	31.7	35.6	17.2	19.7	40.0	19.7	18.1	18.6	5.6	43.2	23.3	19.0	44.1	35.7	27.6
	DIMVC (Xu et al. 2022)	35.5	36.4	18.1	38.7	56.7	17.1	26.7	32.4	12.5	48.2	22.8	19.4	54.0	43.3	33.8
	DCP (Lin et al. 2023)	41.1	45.1	24.8	51.3	74.8	51.9	26.2	32.7	13.5	36.2	18.9	4.8	53.3	42.4	28.6
	SURE (Yang et al. 2023)	41.0	43.2	25.0	43.8	70.1	29.5	25.1	28.3	10.9	52.1	36.9	26.6	57.4	44.8	<u>38.3</u>
	ProImp (Li et al. 2023)	43.6	45.0	26.8	37.6	67.0	25.0	23.7	27.9	10.8	56.5	39.4	32.8	52.2	43.6	31.2
	DIVIDE (Lu et al. 2024)	<u>49.1</u>	<u>48.7</u>	<u>31.6</u>	62.2	83.0	<u>50.5</u>	<u>32.3</u>	<u>39.7</u>	<u>18.1</u>	<u>59.3</u>	<u>39.5</u>	29.0	45.1	30.9	19.4
	CANDY (Guo et al. 2024)	42.0	41.6	24.7	<u>67.3</u>	83.8	60.0	30.6	36.5	16.2	57.7	30.8	<u>37.1</u>	<u>62.1</u>	<u>49.0</u>	37.0
	PMIMC (Yuan et al. 2025)	32.8	36.8	18.4	45.0	71.1	35.4	25.7	33.7	11.6	51.3	31.6	25.3	39.9	34.7	21.6
	Ours	49.7	49.6	32.6	67.9	<u>83.1</u>	49.8	32.6	40.3	18.7	65.3	43.8	39.5	65.9	52.6	47.0
20%	DCCAE (Wang et al. 2015)	30.0	28.3	13.6	45.6	71.1	32.2	21.2	21.6	7.3	42.8	18.4	15.0	43.9	31.1	22.5
	DSIMVC (Tang and Liu 2022)	27.4	27.1	14.1	18.7	29.7	12.4	15.1	14.5	3.7	36.8	16.4	14.0	18.5	13.6	7.2
	DIMVC (Xu et al. 2022)	30.4	30.2	14.1	40.9	56.9	35.3	20.6	25.4	7.4	48.1	23.4	18.4	41.8	30.9	22.4
	DCP (Lin et al. 2023)	38.5	40.2	22.8	44.0	66.3	58.9	26.0	28.4	12.1	38.8	21.6	7.8	47.4	41.2	25.8
	SURE (Yang et al. 2023)	40.1	38.8	22.0	52.1	74.8	39.1	27.3	31.0	13.8	53.3	36.9	28.5	54.3	41.9	35.5
	ProImp (Li et al. 2023)	39.8	<u>42.4</u>	<u>24.2</u>	39.9	68.0	27.3	22.4	26.3	9.8	54.4	<u>39.3</u>	30.8	55.8	44.0	<u>38.0</u>
	DIVIDE (Lu et al. 2024)	<u>41.1</u>	41.0	24.1	67.6	83.1	60.9	31.4	<u>37.7</u>	<u>16.9</u>	<u>57.0</u>	<u>37.6</u>	31.8	59.3	<u>44.6</u>	37.7
	CANDY (Guo et al. 2024)	40.7	38.1	21.6	68.6	83.8	<u>62.7</u>	30.7	33.3	15.1	<u>60.3</u>	38.0	<u>33.9</u>	<u>60.0</u>	<u>37.5</u>	33.6
	PMIMC (Yuan et al. 2025)	33.3	30.9	17.3	<u>37.9</u>	<u>66.2</u>	30.8	18.4	23.3	6.2	<u>46.3</u>	24.3	19.4	26.4	15.0	7.5
	Ours	45.2	42.9	26.9	69.3	83.9	64.4	31.7	38.8	18.3	65.1	42.0	39.4	62.6	47.4	40.8
50%	DCCAE (Wang et al. 2015)	19.6	17.5	6.5	35.0	65.6	27.7	13.4	12.2	2.6	33.1	9.9	12.0	33.8	23.7	14.6
	DSIMVC (Tang and Liu 2022)	26.9	26.1	13.8	16.0	24.0	4.8	11.5	9.1	2.6	34.6	14.3	12.6	12.2	1.5	0.1
	DIMVC (Xu et al. 2022)	23.2	21.9	8.9	27.9	47.5	15.3	13.6	13.6	2.7	44.6	18.7	15.4	33.3	20.1	14.2
	DCP (Lin et al. 2023)	31.4	31.3	16.8	48.3	70.4	59.6	21.5	22.3	8.4	37.0	21.6	5.8	17.7	13.5	8.4
	SURE (Yang et al. 2023)	35.4	<u>34.7</u>	18.2	52.2	75.1	40.9	26.1	28.1	12.6	52.2	36.0	28.7	52.2	39.0	<u>32.7</u>
	ProImp (Li et al. 2023)	<u>36.8</u>	34.4	<u>19.2</u>	34.7	63.7	24.5	19.0	18.2	5.8	41.7	16.7	14.9	50.2	33.9	25.2
	DIVIDE (Lu et al. 2024)	32.5	33.9	17.5	64.3	<u>83.2</u>	59.9	24.0	<u>27.1</u>	9.5	54.8	36.0	<u>31.3</u>	<u>54.7</u>	40.3	31.1
	CANDY (Guo et al. 2024)	34.9	32.7	16.9	<u>67.7</u>	<u>83.2</u>	<u>62.7</u>	23.8	26.3	10.1	<u>54.9</u>	<u>38.4</u>	30.0	48.2	33.9	25.7
	PMIMC (Yuan et al. 2025)	30.7	29.1	15.2	41.2	68.5	31.7	18.0	21.9	5.5	43.2	17.8	17.3	21.3	8.2	4.8
	Ours	42.1	35.3	21.5	73.6	85.9	73.9	<u>25.9</u>	24.4	<u>10.8</u>	61.7	36.6	34.6	55.9	<u>39.2</u>	32.8
80%	DCCAE (Wang et al. 2015)	16.6	15.6	5.1	26.5	58.2	24.0	15.7	13.6	3.5	29.5	8.5	7.4	20.7	14.0	6.8
	DSIMVC (Tang and Liu 2022)	18.9	16.6	9.6	16.1	23.5	5.3	11.8	7.8	1.5	32.5	9.7	9.1	12.6	1.2	0.1
	DIMVC (Xu et al. 2022)	14.9	9.4	2.3	30.7	53.0	27.5	14.5	13.1	2.7	38.7	13.0	8.3	17.8	6.2	1.9
	DCP (Lin et al. 2023)	20.1	20.0	7.2	39.7	61.3	<u>50.4</u>	17.1	16.0	3.2	39.5	25.5	14.9	12.0	1.0	0.1
	SURE (Yang et al. 2023)	36.0	37.7	20.1	50.4	72.8	41.8	25.9	27.8	12.1	46.3	31.2	23.5	40.3	32.2	22.5
	ProImp (Li et al. 2023)	25.8	21.4	10.2	34.3	63.7	25.4	15.5	14.4	4.0	26.6	5.5	4.4	31.8	19.0	12.4
	DIVIDE (Lu et al. 2024)	27.6	29.3	15.5	<u>61.8</u>	<u>80.5</u>	49.2	28.3	34.1	<u>13.6</u>	<u>53.1</u>	34.3	30.2	<u>48.7</u>	<u>33.5</u>	<u>25.2</u>
	CANDY (Guo et al. 2024)	26.5	<u>29.9</u>	14.1	61.5	80.1	48.6	<u>29.0</u>	<u>35.3</u>	13.5	47.2	25.8	22.9	44.0	24.4	20.3
	PMIMC (Yuan et al. 2025)	27.6	23.4	12.0	39.6	67.6	31.4	15.9	19.4	4.9	44.5	17.6	15.2	20.4	6.8	3.7
	Ours	<u>32.4</u>	27.5	<u>16.8</u>	67.8	82.9	57.1	29.2	36.1	15.9	55.5	<u>31.7</u>	<u>27.6</u>	50.6	35.0	30.8

Table 1: The clustering performance on five multi-view datasets at different noise ratios (The missing noise is set to the same ratio as the observation noise)




Figure 2: Ablation study on five dataset with 50% noise ratio

- **Reuters** (Amini, Usunier, and Goutte 2009) is a repository of news content in multiple languages with 18,758 samples. Following (Huang et al. 2019), we transform the texts into a 10-dimensional latent space with a conventional autoencoder and use English and French as two different views.
- **NUS-WIDE** (Hu et al. 2019) includes 9,000 images paired with their respective captions from 10 classes. We adopt a VGG19 neural network for the extraction of visual features, and a Sentence CNN to extract the text features by following (Zhen et al. 2019).

There are three metrics are widely used in MVC, including clustering accuracy (ACC), normalized mutual information (NMI) and adjusted rand index (ARI). Higher values of these indicate better clustering performance.

Implementation Details

The encoder is a four-layer MLP, and the decoder adopts a symmetric architecture per view. The latent feature dimension d is set to 128, and the clustering layer is $d - K$. We use the Adam optimizer with a learning rate of 2×10^{-3} and a batch size of 1024. Temperature parameters τ_c and τ_d are both set to 0.5, the scaling factor σ is 0.07, and the imputation weight α is 0.5. Unless otherwise noted, missing noise ratio η_m and observation noise ratio η_n are both set to 0.5. All experiments are repeated multiple times, and average results are reported. Training is performed on a desktop with an Intel i7-12700KF CPU, NVIDIA RTX 3080Ti GPU, and 32GB RAM using the PyTorch framework. For all baselines, we use publicly available source code and official configurations to ensure fair comparisons.

Comparisons with State of the Arts

We compared the proposed method with nine representative DMVC methods, including DCCA (Wang et al. 2015), DSIMVC (Tang and Liu 2022), DIMVC (Xu et al. 2022), DCP (Lin et al. 2023), SURE (Yang et al. 2023), ProImp (Li et al. 2023), DIVIDE (Lu et al. 2024), CANDY (Guo et al. 2024) and PMIMC (Yuan et al. 2025). As shown in Table 1, our method maintains stable performance in both the noise-free setting and under varying ratios of multi-source noise. These results indicate the effectiveness and robustness of the framework in handling missing noise and observation noise. It is worth noting that DIVIDE and CANDY reduce the impact of false negative pairs by adjusting the contrastive weights by introducing high-order graph structures. However, high-order graphs may amplify noise through propagation. In addition, the presence of false positive pairs can further impair the clustering performance. In contrast, our method leverages a reliability-aware graph to guide contrastive learning under noise, explicitly accounting for the effects of both false negative and false positive pairs. This design effectively enhances the discriminative ability of the learned features and improves the robustness of the model in noisy environments.




Figure 3: Visualization of the training process on Caltech101. 0 and 99 represent the epoch of training.

	Scene15			Caltech101		
	ACC	NMI	ARI	ACC	NMI	ARI
Con	40.8	34.1	20.4	58.7	81.8	39.9
FN con	41.7	34.4	21.1	65.7	82.8	51.1
Ours	42.1	35.3	21.5	73.6	85.9	73.9

Table 2: Effectiveness of noise contrastive learning

Visualization Analysis

To illustrate the data distribution under noisy conditions, we use t -SNE to visualize and compare the training process of the Caltech101 dataset under clean views and 50% multi-source noise. As shown in Fig. 3, before training, the distribution of noisy views is noticeably more scattered than that of the clean views. After 100 epochs of training, the distributions become more aligned. Under noisy conditions, the learned representations have more compact intra-cluster structure and clearer inter-cluster separation compared to the initial state, demonstrating that the our method maintains good clustering performance even in multi-source noise scenarios.

Ablation Study

Effectiveness of design modules: To assess the contribution of each component to the overall performance, we conduct a systematic ablation study. As illustrated in the Fig. 2, "(w/o) dist", "(w/o) crec" and "(w/o) ncon", and "ours" denote the model variants obtained by removing the self-supervised distillation module, cross-view reconstruction module, noise contrastive learning module, and the complete model, respectively. We evaluate all variants on five benchmark datasets using three standard clustering metrics under 50% noise ratio. The results demonstrate that each module contributes positively to performance improvement. In particular, the noise contrastive learning module plays a critical role by enhancing multi-view feature consistency and significantly mitigating the influence of observation noise. Additionally, the cross-view reconstruction mechanism improves robustness by suppressing noise at the data level, while the self-supervised distillation module further refines feature representations and improves clustering performance.

Effectiveness of noise contrastive learning: To analyze the effectiveness of the proposed reliability-aware noise con-

	Reuters			NUS-WIDE		
	ACC	NMI	ARI	ACC	NMI	ARI
Directly	39.9	21.1	13.3	34.4	23.8	9.4
Prototype	51.8	26.7	22.6	52.0	35.3	29.2
k -nn	57.6	32.6	32.1	43.0	31.0	13.9
Ours	61.7	36.6	34.6	55.9	39.2	32.8

Table 3: Effectiveness of imputation




Figure 4: The impact of σ on performance.

trastive learning, we compare it against two baselines: the traditional contrastive learning method (Con) and the false negative pair-aware contrastive method (FN Con) on two datasets under 50% noise ratio. As shown in Table 2, both the traditional and FN contrastive methods fail to perform effectively in noisy environments. In contrast, our method dynamically adjusts the weights of sample pairs based on their estimated reliability. This strategy not only mitigates the effects of both false negative and false positive pairs, but also reduces the impact of noise on the constructions, thereby enhancing the robustness and discriminability of the learned representations.

Effectiveness of imputation: To evaluate the effectiveness of the proposed missing view imputation method, we compare it with three commonly used imputation strategies on two benchmark datasets under 50% noise ratio: direct imputation, prototype imputation, and k -nn imputation. As shown in Table 3, direct imputation replaces the missing view with features from the corresponding view, which may lead to feature homogenization and degradation of multi-view representation quality. Prototype imputation uses the cluster centroid of the observable view for imputation, which introduces discriminative information but compromises view-specific information. k -nn imputation fills in missing data using neighboring features from the same view, preserving view-specific information but ignoring cross-view interactions. In contrast, our method effectively integrates information from multiple views while preserving view-specific features, leading to improved imputation quality and clustering performance.

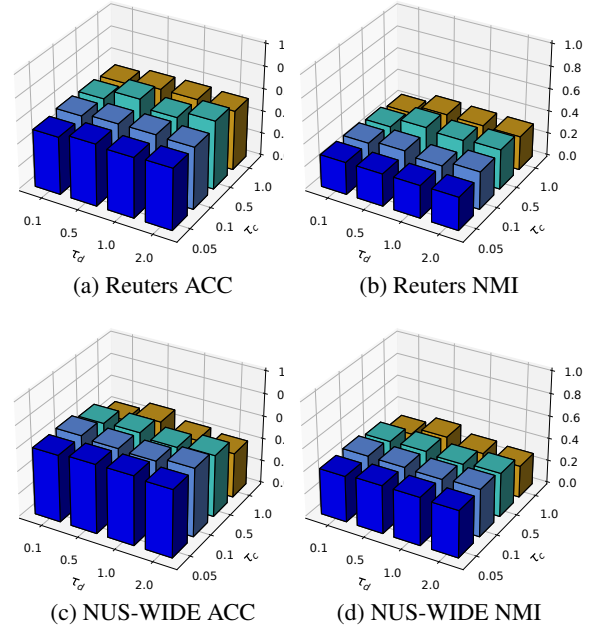


Figure 5: The impact of τ_c, τ_d on performance.

Hyper-parameter Analysis

Sensitivity analysis of hyperparameter σ : To evaluate the effect of the graph construction scaling factor σ on model performance, we conduct experiments on five datasets under 50% noise ratio with $\sigma \in \{0.01, 0.05, 0.1, 0.5, 1.0\}$. As shown in Fig. 4, the model maintains generally stable performance across different values. Notably, performance improves consistently when $\sigma \leq 0.1$, indicating that a smaller σ facilitates better separation of positive and negative pairs and reduces the impact of noisy samples.

Sensitivity analysis of hyperparameter τ_c, τ_d : We further analyzed the impact of the contrastive temperature parameter τ_c and the distillation temperature parameter τ_d on the performance on the Reuters and NUS-WIDE datasets under 50% noise ratio. As shown in Fig. 5, on the Reuters data, when $\tau_c \in [0.05, 0.5]$ and $\tau_d \in [0.5, 1.0]$, the model performs best; on NUS-WIDE, the better performance corresponds to $\tau_c \in [0.05, 0.1]$ and $\tau_d \in [0.1, 0.5]$. This shows that smaller contrastive and distillation temperatures help sharpen the similarity between feature distributions and the discriminability of cluster-level distributions, thereby improving clustering performance.

Conclusion

In this paper, we proposed a reliability-aware contrastive framework for robust multi-view clustering under multi-source noise. The framework is designed to handle both observation noise and view missingness through integrated mechanisms. Extensive experiments on five benchmarks confirm that our method achieves superior performance and maintains robustness across varying noise levels, demonstrating its effectiveness in noisy real-world scenarios.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under grant 92470202 and Frontier Technologies R&D Program of Jiangsu under grant BF2024070.

References

- Abavisani, M.; and Patel, V. M. 2018. Deep multimodal subspace clustering networks. *IEEE Journal of Selected Topics in Signal Processing*, 12(6): 1601–1614.
- Amini, M.; Usunier, N.; and Goutte, C. 2009. Learning from Multiple Partially Observed Views - an Application to Multilingual Text Categorization. In *Advances in Neural Information Processing Systems 22, Vancouver, British Columbia, Canada*, 28–36. Curran Associates, Inc.
- Asad, M.; Azeem, W.; Malik, A. A.; Jiang, H.; Ali, A.; Yang, J.; and Liu, W. 2025. 3D-MMFN: Multi-level multimodal fusion network for 3D industrial image anomaly detection. *Adv. Eng. Informatics*, 65: 103284.
- Cao, X.; Zhang, C.; Fu, H.; Liu, S.; and Zhang, H. 2015. Diversity-induced multi-view subspace clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–594.
- Chao, G.; Jiang, Y.; and Chu, D. 2024. Incomplete Contrastive Multi-View Clustering with High-Confidence Guiding. In *Thirty-Eighth Conference on Artificial Intelligence, Vancouver, Canada*, 11221–11229. AAAI Press.
- Chen, R.; Chen, J.; and Gan, X. 2024. Multi-view graph contrastive learning for social recommendation. *Scientific reports*, 14(1): 22643.
- Cheng, J.; Wang, Q.; Tao, Z.; Xie, D.; and Gao, Q. 2021. Multi-view attribute graph convolution networks for clustering. In *Proceedings of the twenty-ninth international conference on international joint conferences on artificial intelligence*, 2973–2979.
- Cui, H.; Tejada-Lapuerta, A.; Brbić, M.; Saez-Rodriguez, J.; Cristea, S.; Goodarzi, H.; Lotfollahi, M.; Theis, F. J.; and Wang, B. 2025. Towards multimodal foundation models in molecular cell biology. *Nature*, 640(8059): 623–633.
- Fei-Fei, L.; and Perona, P. 2005. A Bayesian Hierarchical Model for Learning Natural Scene Categories. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, San Diego, CA*, 524–531. IEEE Computer Society.
- Feng, W.; Sheng, G.; Wang, Q.; Gao, Q.; Tao, Z.; and Dong, B. 2024. Partial Multi-View Clustering via Self-Supervised Network. In *Thirty-Eighth Conference on Artificial Intelligence, Vancouver, Canada*, 11988–11995. AAAI Press.
- Gao, H.; Nie, F.; Li, X.; and Huang, H. 2015. Multi-view subspace clustering. In *Proceedings of the IEEE international conference on computer vision*, 4238–4246.
- Guo, R.; Yang, M.; Lin, Y.; Peng, X.; and Hu, P. 2024. Robust Contrastive Multi-view Clustering against Dual Noisy Correspondence. In *Advances in Neural Information Processing Systems 38, Vancouver, BC, Canada*.
- Han, Z.; Zhang, C.; Fu, H.; and Zhou, J. T. 2021. Trusted Multi-View Classification. In *9th International Conference on Learning Representations, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Holm, N. N.; Le, T. M.; Frølich, A.; Andersen, O.; Juul-Larsen, H. G.; Stockmarr, A.; and Venkatesh, S. 2025. am-VAE: Age-aware multimorbidity clustering using variational autoencoders. *Computers in Biology and Medicine*, 186: 109632.
- Hu, P.; Zhen, L.; Peng, D.; and Liu, P. 2019. Scalable Deep Multimodal Learning for Cross-Modal Retrieval. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, Paris, France*, 635–644. ACM.
- Huang, F.; Zhang, X.; Li, C.; Li, Z.; He, Y.; and Zhao, Z. 2018. Multimodal Network Embedding via Attention based Multi-view Variational Autoencoder. In *Proceedings of the 2018 ACM on International Conference on Multimedia Retrieval, Yokohama, Japan*, 108–116. ACM.
- Huang, S.; Kang, Z.; and Xu, Z. 2020. Auto-weighted multi-view clustering via deep matrix decomposition. *Pattern Recognition*, 97: 107015.
- Huang, Z.; Zhou, J. T.; Peng, X.; Zhang, C.; Zhu, H.; and Lv, J. 2019. Multi-view Spectral Clustering Network. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, Macao*, 2563–2569. ijcai.org.
- Jin, J.; Wang, S.; Dong, Z.; Liu, X.; and Zhu, E. 2023. Deep Incomplete Multi-View Clustering with Cross-View Partial Sample and Prototype Alignment. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, Vancouver, BC, Canada*, 11600–11609. IEEE.
- Kang, Z.; Zhou, W.; Zhao, Z.; Shao, J.; Han, M.; and Xu, Z. 2020. Large-scale multi-view subspace clustering in linear time. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 4412–4419.
- Li, H.; Li, Y.; Yang, M.; Hu, P.; Peng, D.; and Peng, X. 2023. Incomplete Multi-view Clustering via Prototype-based Imputation. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, Macao, SAR*, 3911–3919. ijcai.org.
- Li, J.; and Zhou, X. 2025. CureGraph: Contrastive multimodal graph representation learning for urban living circle health profiling and prediction. *Artificial Intelligence*, 340: 104278.
- Li, Y.; Nie, F.; Huang, H.; and Huang, J. 2015. Large-Scale Multi-View Spectral Clustering via Bipartite Graph. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, Austin, Texas*, 2750–2756. AAAI Press.
- Li, Z.; Wang, Q.; Tao, Z.; Gao, Q.; Yang, Z.; et al. 2019. Deep Adversarial Multi-view Clustering Network. In *IJCAI*, volume 2, 4.
- Lin, Y.; Gou, Y.; Liu, X.; Bai, J.; Lv, J.; and Peng, X. 2023. Dual Contrastive Prediction for Incomplete Multi-View Representation Learning. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(4): 4447–4461.

- Liu, J.; Wang, C.; Gao, J.; and Han, J. 2013. Multi-view clustering via joint nonnegative matrix factorization. In *Proceedings of the 2013 SIAM international conference on data mining*, 252–260.
- Lu, Y.; Lin, Y.; Yang, M.; Peng, D.; Hu, P.; and Peng, X. 2024. Decoupled Contrastive Multi-View Clustering with High-Order Random Walks. In *Thirty-Eighth Conference on Artificial Intelligence, Vancouver, Canada*, 14193–14201. AAAI Press.
- Pu, J.; Cui, C.; Chen, X.; Ren, Y.; Pu, X.; Hao, Z.; Yu, P. S.; and He, L. 2024. Adaptive Feature Imputation with Latent Graph for Deep Incomplete Multi-View Clustering. In *Thirty-Eighth Conference on Artificial Intelligence, February 20-27, 2024, Vancouver, Canada*, 14633–14641. AAAI Press.
- Rao, V. M.; Hla, M.; Moor, M.; Adithan, S.; Kwak, S.; Topol, E. J.; and Rajpurkar, P. 2025. Multimodal generative AI for medical image interpretation. *Nature*, 639(8056): 888–896.
- Ren, Y.; Pu, J.; Cui, C.; Zheng, Y.; Chen, X.; Pu, X.; and He, L. 2024. Dynamic Weighted Graph Fusion for Deep Multi-View Clustering. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, 4842–4850. ijcai.org.
- Sun, Y.; Li, Y.; Ren, Z.; Duan, G.; Peng, D.; and Hu, P. 2025. ROLL: Robust Noisy Pseudo-label Learning for Multi-View Clustering with Noisy Correspondence. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, Nashville, TN, USA*, 30732–30741. Computer Vision Foundation / IEEE.
- Tang, C.; Liu, X.; Zhu, X.; Zhu, E.; Luo, Z.; Wang, L.; and Gao, W. 2020. CGD: Multi-view clustering via cross-view graph diffusion. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 5924–5931.
- Tang, H.; and Liu, Y. 2022. Deep Safe Incomplete Multi-view Clustering: Theorem and Algorithm. In *International Conference on Machine Learning, Baltimore, Maryland*, volume 162 of *Proceedings of Machine Learning Research*, 21090–21110. PMLR.
- Tang, X.; Tang, X.; Wang, W.; Fang, L.; and Wei, X. 2018. Deep Multi-view Sparse Subspace Clustering. In *Proceedings of the VII International Conference on Network, Communication and Computing, Taipei City, Taiwan*, 115–119.
- Wang, H.; Yang, Y.; and Liu, B. 2019. GMC: Graph-based multi-view clustering. *IEEE Transactions on Knowledge and Data Engineering*, 32(6): 1116–1129.
- Wang, W.; Arora, R.; Livescu, K.; and Bilmes, J. A. 2015. On Deep Multi-View Representation Learning. In *Proceedings of the 32nd International Conference on Machine Learning, Lille, France*, volume 37 of *JMLR Workshop and Conference Proceedings*, 1083–1092. JMLR.org.
- Wang, Y.; Peng, J.; Zhang, J.; Yi, R.; Wang, Y.; and Wang, C. 2023. Multimodal Industrial Anomaly Detection via Hybrid Fusion. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, Vancouver, BC, Canada*, 8032–8041. IEEE.
- Wen, J.; Deng, S.; Wong, W.; Chao, G.; Huang, C.; Fei, L.; and Xu, Y. 2024. Diffusion-based Missing-view Generation With the Application on Incomplete Multi-view Clustering. In *Forty-first International Conference on Machine Learning, Vienna, Austria*. OpenReview.net.
- Xu, C.; Guan, Z.; Zhao, W.; Wu, H.; Niu, Y.; and Ling, B. 2019. Adversarial incomplete multi-view clustering. In *IJ-CAI*, volume 7, 3933–3939.
- Xu, J.; Li, C.; Ren, Y.; Peng, L.; Mo, Y.; Shi, X.; and Zhu, X. 2022. Deep Incomplete Multi-View Clustering via Mining Cluster Complementarity. In *Thirty-Sixth Conference on Artificial Intelligence*, 8761–8769. AAAI Press.
- Xu, J.; Ren, Y.; Tang, H.; Pu, X.; Zhu, X.; Zeng, M.; and He, L. 2021. Multi-VAE: Learning disentangled view-common and view-peculiar visual representations for multi-view clustering. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9234–9243.
- Xu, J.; Ren, Y.; Wang, X.; Feng, L.; Zhang, Z.; Niu, G.; and Zhu, X. 2024. Investigating and Mitigating the Side Effects of Noisy Views for Self-Supervised Clustering Algorithms in Practical Multi-View Scenarios. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA*, 22957–22966. IEEE.
- Yang, M.; Li, Y.; Hu, P.; Bai, J.; Lv, J.; and Peng, X. 2023. Robust Multi-View Clustering With Incomplete Information. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(1): 1055–1069.
- Yang, X.; Wang, S.; Wang, F.; Jin, J.; Liu, S.; Liu, Y.; Zhu, E.; Liu, X.; and Jin, Y. 2025. Automatically Identify and Rectify: Robust Deep Contrastive Multi-view Clustering in Noisy Scenarios. *arXiv preprint arXiv:2505.21387*.
- Yang, Y.; and Newsam, S. D. 2010. Bag-of-visual-words and spatial extensions for land-use classification. In *18th International Symposium on Advances in Geographic Information Systems, San Jose, CA, Proceedings*, 270–279. ACM.
- Yuan, H.; Sun, Y.; Zhou, F.; Wen, J.; Yuan, S.; You, X.; and Ren, Z. 2025. Prototype Matching Learning for Incomplete Multi-View Clustering. *IEEE Trans. Image Process.*, 34: 828–841.
- Zhang, Y.; Lin, Y.; Yan, W.; Yao, L.; Wan, X.; Li, G.; Zhang, C.; Ke, G.; and Xu, J. 2025a. Incomplete Multi-view Clustering via Diffusion Contrastive Generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 22650–22658.
- Zhang, Y.; Yan, W.; Tang, C.; Zhou, W.; and Jin, J. 2025b. Multi-branch Space Sharing Feature Aggregation for contrastive multi-view clustering. *Pattern Recognition*, 111704.
- Zhao, H.; Ding, Z.; and Fu, Y. 2017. Multi-view clustering via deep matrix factorization. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31.
- Zhen, L.; Hu, P.; Wang, X.; and Peng, D. 2019. Deep Supervised Cross-Modal Retrieval. In *IEEE Conference on Computer Vision and Pattern Recognition, Long Beach, CA*, 10394–10403. Computer Vision Foundation / IEEE.