

Deep Incomplete Multi-View Clustering via Hierarchical Imputation and Alignment

Yiming Du, Ziyu Wang, Jian Li, Rui Ning, Lusi Li*

Department of Computer Science, Old Dominion University, Norfolk, VA 23529, USA
 {ydu002, zwang007, jli038}@odu.edu, {rning, lusili}@cs.odu.edu

Abstract

Incomplete multi-view clustering (IMVC) aims to discover shared cluster structures from multi-view data with partial observations. The core challenges lie in accurately imputing missing views without introducing bias, while maintaining semantic consistency across views and compactness within clusters. To address these challenges, we propose DIMVC-HIA, a novel deep IMVC framework that integrates hierarchical imputation and alignment with four key components: (1) view-specific autoencoders for latent feature extraction, coupled with a view-shared clustering predictor to produce soft cluster assignments; (2) a hierarchical imputation module that first estimates missing cluster assignments based on cross-view contrastive similarity, and then reconstructs missing features using intra-view, intra-cluster statistics; (3) an energy-based semantic alignment module, which promotes intra-cluster compactness by minimizing energy variance around low-energy cluster anchors; and (4) a contrastive assignment alignment module, which enhances cross-view consistency and encourages confident, well-separated cluster predictions. Experiments on benchmarks demonstrate that our framework achieves superior performance under varying levels of missingness.

Code — <https://github.com/YMBest/DIMVC-HIA>

Introduction

Multi-view clustering (MVC) has emerged as a fundamental paradigm in unsupervised learning, with broad applications in multimedia analysis (Cao et al. 2025), bioinformatics (Zang et al. 2025), and social network mining (Zeng, Peng, and Li 2024). In real-world data representation, information is often collected from multiple heterogeneous sources—such as different sensors, modalities, or feature extractors—resulting in multi-view data that provides complementary representations of the same underlying phenomena. By integrating these diverse perspectives, MVC offers richer insights than single-view approaches (Lu et al. 2024; Guan et al. 2025), enabling more robust pattern discov-

ery, improved generalization, and a deeper understanding of complex data structures.

Despite these advantages, conventional MVC methods rely on a critical yet often unrealistic assumption that all views are complete and fully observable (Chao, Jiang, and Chu 2024). Missing views are pervasive in real-world scenarios, arising from sensor failures, data corruption, transmission errors, or preprocessing artifacts. This incompleteness severely limits the applicability of MVC techniques. To address this gap, incomplete MVC (IMVC) has emerged as an active research frontier, aiming to perform reliable clustering when some views are partially missing—sometimes leaving only one available view. The primary challenge of IMVC is to leverage available information effectively while mitigating the adverse effects of missing data, ensuring that clustering reliability and accuracy remain intact.

Recent advancements in IMVC have led to the development of various methodologies to handle missing views, including graph-based (Tan et al. 2024; Lu et al. 2024), matrix factorization-based (Park and Lee 2025), kernel-based (Liu et al. 2023; Ding and Yang 2024), subspace-based (Ji and Feng 2025; Gu et al. 2024), and deep learning-based techniques (Du et al. 2025; Wang et al. 2025b; Dong et al. 2025). These methods generally fall into two categories. Imputation-based approaches (Pu et al. 2024; Liu et al. 2024a; Tu et al. 2024) attempt to reconstruct missing views through data-level completion, feature-space recovery, or consensus-driven imputation by leveraging cross-view dependencies. While potentially effective when reconstructions are accurate, these methods suffer from a critical vulnerability: error propagation. Poor imputations distort structural patterns, which in turn degrade subsequent imputations, creating a self-reinforcing cycle of deteriorating clustering performance.

On the other hand, imputation-free approaches (Liu et al. 2024b; Dai et al. 2025) avoid explicit data recovery by learning shared latent spaces directly from available views or performing clustering without view reconstruction. These methods demonstrate greater robustness under moderate missingness but face escalating challenges as missing data increases: (i) instance misalignment, where missing views disrupt sample correspondences, hindering the construction of a coherent latent space (Wang et al. 2025c); (ii) information imbalance, where dominant views overshadow less-

*Corresponding Author

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

represented ones, distorting feature learning and diminishing the contribution of incomplete views (Xu et al. 2024); and (iii) uncertainty in representation learning and clustering assignments, where high missingness weakens confidence in sample-to-cluster assignments (Chen et al. 2025). This uncertainty is particularly critical for instances positioned near cluster boundaries or within closely related but distinct clusters, where small variations in representation can lead to misclassification during clustering. The resulting unreliable clustering decisions exacerbate inaccuracies in both learned features and the identification of latent patterns.

To address these challenges, we propose Deep Incomplete Multi-View Clustering via Hierarchical Imputation and Alignment (DIMVC-HIA), which integrates view-specific representation learning with a two-stage hierarchical imputation strategy and dual alignment mechanisms. As illustrated in Figure 1, our framework first employs view-specific autoencoders to learn latent representations coupled with a shared predictor for soft clustering assignments, then performs hierarchical imputation by (i) recovering missing cluster assignment distributions through cross-view contrastive similarity to establish semantic priors, followed by (ii) structure-preserving latent feature reconstruction using intra-view, intra-cluster statistics. To ensure robust clustering, we introduce dual alignment: energy-based models (EBMs) (Bachtis et al. 2024; Peng et al. 2024) minimize intra-cluster energy variance to enhance semantic coherence, while a contrastive objective aligns cluster assignments across views to maintain inter-view consistency and promote well-separated clusters. This unified architecture addresses the core challenges of error propagation in imputation-based methods and representation uncertainty in imputation-free approaches, achieving both accurate data recovery and reliable clustering. Our contributions can be summarized as follows:

- We propose DIMVC-HIA, a novel DIMVC framework that jointly optimizes view-specific representation learning, hierarchical imputation, and dual-alignment regularization within an end-to-end architecture, effectively bridging the gap between imputation-based and imputation-free paradigms.
- We develop a hierarchical imputation strategy that first recovers missing cluster assignments using inter-view contrastive similarity, and then reconstructs missing features conditioned on these imputed cluster structures using intra-view cluster statistics, ensuring both geometric and semantic consistency.
- We design dual alignment modules: (i) an energy-based alignment that promotes compact clusters using cluster-specific energy models, and (ii) a contrastive assignment alignment that improves inter-view consistency and encourages confident, well-separated predictions.

Related Work

Deep Incomplete Multi-View Clustering

DIMVC leverages the powerful representation learning capabilities of deep neural networks such as autoencoders

(AEs), generative adversarial networks (GANs), and graph neural networks (GNNs) to reduce dimensionality and learn shared latent structures across views, enabling effective handling of incomplete or missing data. Among various strategies, contrastive learning (Shiri, Reddy, and Sun 2024; Sun et al. 2024) has emerged as a prominent and widely adopted self-supervised learning approach. Its fundamental objective is to maximize the agreement between positive pairs while minimizing the similarity between negative pairs (Dou et al. 2025). This can be applied at the instance level, cluster level, or both, to enhance feature representation and improve clustering performance. For instance, Wang et al. (Wang et al. 2025a) employed cross-view contrastive alignments at both the instance and cluster levels, effectively capturing consistent assignments. Xue et al. (Xue et al. 2022) proposed a diversified graph contrastive regularization strategy at the intra-graph, inter-graph, and clustering levels to capture diverse data correlations, enhance discriminative representation learning, and mitigate information loss.

In addition, many DIMVC methods adopt imputation-based strategies (Yang et al. 2025; Kim, Lee, and Park 2025) to recover missing views and enable more complete representation learning. Some approaches infer missing data by exploiting structural relationships across views, such as sample similarity, graph connectivity, or prototype semantics, often using tools like GNNs or clustering-based alignment (Li, Zhao, and Qin 2021; Zhang, Wang, and Yang 2023). Others adopt predictive or generative models to reconstruct missing features from observed ones (Li, Xu, and Lu 2019). Representative methods such as MICA (Wang et al. 2025c) follow this paradigm by inferring assignments from imputed features, making it directly affected by imputation noise. While these methods have shown promise, they face challenges when observed views are noisy or insufficient, potentially leading to inaccurate imputation or distribution mismatch.

Energy-based Models

EBMs offer a principled approach to modeling complex data distributions by associating each data sample with a scalar energy value that reflects its plausibility (LeCun et al. 2006; Salakhutdinov and Hinton 2009). Formally, a probability density $p_\theta(x)$ for $x \in \mathbb{R}^d$ can be defined via the Boltzmann distribution:

$$p_\theta(x) = \frac{\exp(-\mathcal{E}_\theta(x))}{Z_\theta}, \quad (1)$$

where $\mathcal{E}_\theta(x) : \mathbb{R}^d \rightarrow \mathbb{R}$ denotes the parameterized energy function that assigns a scalar energy to each input x , and $Z_\theta = \int \exp(-\mathcal{E}_\theta(x)) dx$ is the partition function that ensures proper normalization. In practice, computing Z_θ is intractable for high-dimensional data, but this is often unnecessary. Many training and inference procedures avoid direct dependence on Z_θ by using energy differences or relative probabilities, which cancel out the partition function (Nie, Vahdat, and Anandkumar 2021; Margeloiu et al. 2024). This allows EBMs to directly shape the energy landscape—assigning lower energy to likely samples and higher energy to unlikely ones—without requiring explicit normalization.

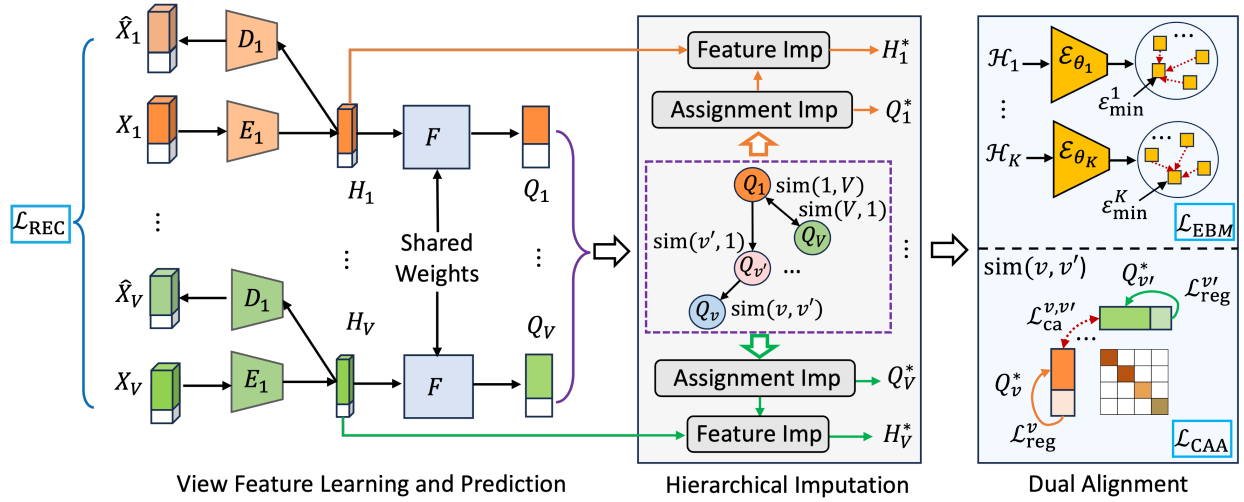


Figure 1: The overall architecture of our proposed DIMVC-HIA framework.

Methodology

Notations. Let $X = \{X_v \in \mathbb{R}^{N \times d_v}\}_{v=1}^V$ denote an incomplete multi-view dataset with V views and up to N samples per view. For each view v , X_v is the data matrix, where the i -th sample is $X_v(i) \in \mathbb{R}^{d_v}$ and its (i, j) -th entry is $X_v(i, j)$. The available samples in view v are denoted as $X_v^a \in \mathbb{R}^{N_v \times d_v}$ with $N_v \leq N$, and missing entries are padded with zeros. A binary indicator matrix $G \in \{0, 1\}^{N \times V}$ marks sample availability, with $G(i, v) = 1$ if $X_v(i)$ is observed, and 0 otherwise. The goal is to uncover K shared cluster structures across views.

View Feature Learning and Prediction

We first adopt view-specific autoencoders to extract features from heterogeneous multi-view data. For data X_v from view v , an encoder E_v and decoder D_v are used to extract latent features $H_v \in \mathbb{R}^{N \times d}$, and reconstruct the input:

$$H_v = E_v(X_v; \phi_v^e), \quad \hat{X}_v = D_v(H_v; \phi_v^d). \quad (2)$$

Here, d is the dimension of the shared latent space, and θ_v , ϕ_v are the parameters of the encoder and decoder, respectively. Each AE is pre-trained independently by minimizing the reconstruction loss over available samples:

$$\mathcal{L}_{\text{REC}} = \frac{1}{VN} \sum_{v=1}^V \sum_{i=1}^N G(i, v) \|X_v(i) - \hat{X}_v(i)\|_2^2. \quad (3)$$

The indicator $G(i, v)$ ensures that only observed samples contribute to the loss.

The latent features H_v are subsequently mapped to soft cluster assignments via a shared clustering predictor F parameterized by ϑ :

$$Q_v(i) = F(H_v(i); \vartheta) \in \mathbb{R}^K, \quad (4)$$

where $Q_v(i)$ denotes the cluster assignment distribution for the i -th sample in view v .

Hierarchical Imputation

We argue that the latent space of soft cluster assignments encodes richer semantic information that more directly reflects the underlying cluster structures. Therefore, we propose to first explore cross-view relationships within this assignment space and leverage them to guide the imputation of missing cluster assignments. The imputed assignments are then used to inform the imputation of latent feature representations.

Cluster Assignment Imputation. To impute missing cluster assignments, we exploit semantic relationships across views by constructing pairwise similarity graphs from co-observed samples. The key idea is that samples observed in multiple views should share consistent clustering semantics. These correspondences enable us to infer missing assignments in one view using reliable predictions from semantically aligned views. For each distinct view pair (v, v') , where $v \neq v'$ and $v, v' \in \{1, \dots, V\}$, we define the index set of co-observed samples as:

$$\mathcal{I}_{v, v'} = \{i \mid G(i, v) = 1 \text{ and } G(i, v') = 1\} \quad (5)$$

Let $|\mathcal{I}_{v, v'}|$ be the number of such paired samples. We then extract the soft cluster assignments for these samples in views v and v' as:

$$Q_v^{v, v'} = [Q_v(j)]_{j \in \mathcal{I}_{v, v'}}, \quad Q_{v'}^{v', v} = [Q_{v'}(j)]_{j \in \mathcal{I}_{v, v'}} \quad (6)$$

where $Q_v^{v, v'}, Q_{v'}^{v', v} \in \mathbb{R}^{|\mathcal{I}_{v, v'}| \times K}$. The cross-view similarity matrix between these two sets of assignments is computed as:

$$S_{v, v'} = Q_v^{v, v'} (Q_{v'}^{v', v})^\top \in \mathbb{R}^{|\mathcal{I}_{v, v'}| \times |\mathcal{I}_{v, v'}|}. \quad (7)$$

To quantify the semantic alignment between two views, we propose a label-aware contrastive similarity score. For each co-observed sample $i \in \mathcal{I}_{v, v'}$ in view v , the positive similarity is defined as the diagonal entry $S_{v, v'}(i, i)$, which measures agreement between the two views for the same instance. The negative similarities are given by off-diagonal entries $S_{v, v'}(i, j)$ where $j \neq i$ and the predicted

cluster labels differ, i.e., $\hat{y}_v(i) \neq \hat{y}_{v'}(j)$. Here, $\hat{y}_v(i) = \arg \max_k Q_v(i, k)$ and $\hat{y}_{v'}(j) = \arg \max_k Q_{v'}(j, k)$.

To avoid penalizing semantically consistent samples, we exclude false negatives—samples $j \neq i$ that share the same predicted label as i . Let $\mathcal{F}_{v,v'}^i = \{j \in \mathcal{I}_{v,v'} \mid j \neq i, \hat{y}_v(i) = \hat{y}_{v'}(j)\}$ denote the set of false negatives in view v' with respect to sample i in view v . The set of valid negative candidates is then given by: $\mathcal{B}_{v,v'}^i = \{\mathcal{I}_{v,v'} \setminus \mathcal{F}_{v,v'}^i\}$. Using these filtered negatives, the contrastive similarity score between views v and v' is computed as:

$$\text{sim}(v, v') = \frac{1}{|\mathcal{I}_{v,v'}|} \sum_{i \in \mathcal{I}_{v,v'}} \frac{\exp(S_{v,v'}(i, i)/\tau)}{\sum_{j \in \mathcal{B}_{v,v'}^i} \exp(S_{v,v'}(i, j)/\tau)}. \quad (8)$$

This label-aware contrastive formulation improves semantic discriminability by highlighting inconsistencies and reducing the impact of noisy similarities. The resulting scores guide assignment imputation by prioritizing more semantically aligned views.

With the computed cross-view similarity scores $\text{sim}(\cdot, \cdot)$, each target view v defines a reference view list \mathcal{R}_v , which contains all other views sorted in descending order of semantic similarity:

$$\mathcal{R}_v = \text{argsort}_{v' \neq v} (\text{sim}(v, v')). \quad (9)$$

For each missing sample i in view v , we sequentially traverse the ranked list \mathcal{R}_v and select the first view $v' \in \mathcal{R}_v$ such that sample i is available in view v' . We denote this selected reference view as:

$$\pi_v^i = \min \{v' \in \mathcal{R}_v \mid G(i, v') = 1\}. \quad (10)$$

The imputed cluster assignment for sample i in view v is then obtained from the most semantically aligned available view π_v^i , and the completed assignment distribution is defined as:

$$Q_v^*(i) = \begin{cases} Q_v(i), & \text{if } G(i, v) = 1 \\ Q_{\pi_v^i}(i), & \text{otherwise.} \end{cases} \quad (11)$$

This similarity-guided imputation strategy ensures that the missing cluster assignment probabilities are inferred from the semantically aligned and structurally reliable views.

Latent Feature Imputation. After completing the cluster assignments Q_v^* for each view v via inter-view similarity, we further leverage this information to impute missing latent features in a cluster-aware intra-view manner.

For each missing sample i in view v , we first determine its most likely cluster label:

$$\hat{y}_v(i) = \arg \max_k Q_v^*(i, k), \quad k \in \{1, \dots, K\}. \quad (12)$$

We then compute the cluster prototype $\mathcal{C}_v(k)$ as the average latent feature of all available samples in view v that are predicted to belong to cluster k :

$$\mathcal{C}_v(k) = \frac{1}{|Y_v^k|} \sum_{j \in Y_v^k} H_v(j) \quad (13)$$

where $Y_v^k = \{j \mid G(j, v) = 1, \hat{y}_v(j) = k\}$. The completed latent feature $H_v^*(i)$ is then obtained by:

$$H_v^*(i) = \begin{cases} H_v(i), & \text{if } G(i, v) = 1 \\ \mathcal{C}_v(\hat{y}_v(i)), & \text{otherwise.} \end{cases} \quad (14)$$

Although non-parametric, these imputations influence learning through dual-alignment losses, while reconstruction on observed samples stabilizes training.

Energy-Based Semantic Alignment

Given the completed features H_v^* and predicted cluster labels \hat{y}_v , we introduce an energy-based semantic alignment module to enhance intra-cluster compactness and semantic consistency across views. This module leverages cluster-aware EBMs to assess the reliability of latent features—assigning lower energy to those more semantically aligned with their clusters. For each cluster $k \in \{1, \dots, K\}$, we define a view-shared, cluster-specific energy function $\mathcal{E}_{\theta_k} : \mathbb{R}^d \mapsto \mathbb{R}^+$, which maps feature vectors to scalar energy scores. A lower score indicates stronger compatibility with cluster k . We first construct the cluster-level feature set by collecting all features across views that are assigned to cluster k :

$$\mathcal{H}_k = \{H_v^*(i) \mid \hat{y}_v(i) = k, \forall v, i\}. \quad (15)$$

To promote semantic alignment within each cluster, we minimize the energy deviation of features from the most confident cluster anchor—the feature with the lowest energy:

$$\varepsilon_{\min}^k = \min_{\mathbf{h}' \in \mathcal{H}_k} \mathcal{E}_{\theta_k}(\mathbf{h}'; \theta_k). \quad (16)$$

The energy-based alignment loss for cluster k is then computed to enforce intra-cluster compactness:

$$\mathcal{L}_{\text{EBM}}^k = \frac{1}{|\mathcal{H}_k|} \sum_{\mathbf{h} \in \mathcal{H}_k} |\mathcal{E}_{\theta_k}(\mathbf{h}; \theta_k) - \varepsilon_{\min}^k|. \quad (17)$$

The overall energy-based alignment loss is obtained by averaging over all clusters:

$$\mathcal{L}_{\text{EBM}} = \frac{1}{K} \sum_{k=1}^K \mathcal{L}_{\text{EBM}}^k. \quad (18)$$

This formulation encourages features within each cluster to concentrate around reliable, low-energy anchors. It enables the model to flexibly shape continuous energy landscapes beyond center-based regularizers.

Contrastive Assignment Alignment

Given the completed soft cluster assignments Q_v^* , we introduce a contrastive assignment alignment (CAA) loss that promotes both sample-level semantic agreement and confident, well-separated cluster predictions. The total loss is defined as:

$$\mathcal{L}_{\text{CAA}} = \frac{1}{2} \sum_{v=1}^V \sum_{v'=1, v' \neq v}^V [\text{sim}(v, v') \cdot \mathcal{L}_{\text{ca}}^{v,v'} + \mathcal{L}_{\text{reg}}^{v,v'}]. \quad (19)$$

Here, $\text{sim}(v, v')$ represents the semantic similarity score between views v and v' , as defined in Eq. (8). The alignment loss comprises two components.

The contrastive alignment loss is used to pull together the cluster assignment distributions of matching samples across views:

$$\mathcal{L}_{ca}^{v,v'} = -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp(Q_v^*(i)^\top Q_{v'}^*(i)/\tau)}{\sum_{j=1}^N \exp(Q_v^*(i)^\top Q_{v'}^*(j)/\tau)}, \quad (20)$$

where τ is a temperature scaling factor. To promote confident and balanced predictions, we incorporate a distribution-level entropy regularization. For each view, we compute the average cluster assignment distribution:

$$q_v(k) = \frac{1}{N} \sum_{i=1}^N Q_v^*(i, k), \quad q_{v'}(k) = \frac{1}{N} \sum_{j=1}^N Q_{v'}^*(j, k). \quad (21)$$

The entropy regularization term is then given by:

$$\mathcal{L}_{reg}^{v,v'} = \frac{1}{K} \sum_{k=1}^K [q_v(k) \log(q_v(k)) + q_{v'}(k) \log(q_{v'}(k))]. \quad (22)$$

The entropy regularizer promotes balanced cluster proportions and prevents trivial all-in-one or empty assignments.

Overall Objective

The overall training objective function is defined as:

$$\mathcal{L} = \mathcal{L}_{REC} + \alpha \cdot \mathcal{L}_{EBM} + \beta \cdot \mathcal{L}_{CAA}. \quad (23)$$

Here, \mathcal{L}_{REC} denotes the reconstruction loss for pretraining the autoencoders, \mathcal{L}_{EBM} encourages intra-cluster compactness across views, and \mathcal{L}_{CAA} enforces cross-view consistency and confidence through contrastive assignment alignment. The hyperparameters α and β balance the contribution of the alignment terms. The full set of trainable model parameters, denoted as Φ , includes the view-specific autoencoder parameters $\{\phi_v^e, \phi_v^d\}_{v=1}^V$, the view-shared clustering predictor parameters ϑ , and the cluster-specific EBM parameters $\{\theta_k\}_{k=1}^K$. The training procedure is summarized in Algorithm 1. After training, the final cluster label for sample i is determined by $y_i = \arg \max_k (\sum_v Q_v^*(i, k))$, and the final labels $Y = [y_1, \dots, y_N]$.

Experiments

Experimental settings

Datasets and Evaluation Metrics We conducted experiments on four benchmark multi-view datasets: BDGP (Cai et al. 2012), MNIST-USPS (Asuncion, Newman et al. 2007), Fashion-MNIST (Xiao, Rasul, and Vollgraf 2017), and Handwritten (Asuncion, Newman et al. 2007), as summarized in Table 1. The incomplete versions of these datasets are generated by randomly omitting a proportion of samples from each view, corresponding to missing ratios $\eta \in \{0.1, 0.3, 0.5, 0.7\}$. Three metrics, ACC, NMI, and Purity (PUR), are adopted to evaluate the clustering performance.

Baseline Methods and Evaluation Metrics We compare the performance of DIMVC-HIA against ten state-of-the-art incomplete multi-view clustering methods. The evaluated baselines include: COMPLETER (Lin et al. 2021),

Algorithm 1 Training Procedure of DIMVC-HIA

- 1: **Input:** Incomplete multi-view data X , missing indicator matrix G , hyperparameters α, β, τ
 - 2: **Output:** Final clustering assignments Y
 - 3: Initialize model parameters Φ
 - 4: Pretrain autoencoders by minimizing \mathcal{L}_{REC} via Eq. (3)
 - 5: **for** $t = 1$ to T **do**
 - 6: Sample a mini-batch $\{X_v(i)\}_{i=1}^B$ from each view X_v
 - 7: Compute features $\{H_v(i)\}_{i=1}^B$ and soft cluster assignments $\{Q_v(i)\}_{i=1}^B$ via Eqs. (2) and (4), respectively
 - 8: Compute inter-view semantic similarity scores $\text{sim}(\cdot, \cdot)$ via Eq. (8)
 - 9: Generate completed cluster assignments $\{Q_v^*(i)\}_{i=1}^B$ for each view via Eq. (11)
 - 10: Generate completed features $\{H_v^*(i)\}_{i=1}^B$ for each view via Eq. (14)
 - 11: Compute \mathcal{L}_{EBM} and \mathcal{L}_{CAA} via Eqs. (18) and (19)
 - 12: Compute the total training loss \mathcal{L} via Eq. (23)
 - 13: Update model parameters Φ using backpropagation
 - 14: **end for**
 - 15: Compute final cluster labels Y
-

Dataset	N / V / K	Dimensions
BDGP	2,500 / 2 / 5	1,750/79
MNIST-USPS	5,000 / 2 / 10	784/256
Fashion-MNIST	10,000 / 3 / 10	784/784/784
Handwritten	2,000 / 6 / 10	240/76/216/47/64/6

N (# of Samples); V (# of Views); K (# of Clusters).

Table 1: Statistics of the benchmark datasets.

CPSPAN (Jin et al. 2023), DCG (Zhang et al. 2025), DIVIDE (Lu et al. 2024), RPCIC (Yuan et al. 2024), APADC (Xu et al. 2023), GIMVC (Bai et al. 2024), PMIMC (Yuan et al. 2025), ProImp (Li et al. 2023), and DSIMVC (Tang and Liu 2022).

Implementation Details DIMVC-HIA employs view-specific encoders with architectures $[d_v, 256, 512, d]$ for BDGP and Fashion-MNIST, and $[d_v, 256, 512, 1024, d]$ for MNIST-USPS and Handwritten, where d_v denotes the input dimension and $d = 2000$ is fixed across all datasets. The shared MLP comprises two linear layers with dimensions $[d, 1024, K]$, where K is the number of clusters. Each cluster-specific EBM is an MLP with hidden dimensions $[256, 256, 256]$. We use batch sizes of 200 for BDGP and Handwritten, 100 for Fashion-MNIST, and 50 for MNIST-USPS, with a fixed learning rate of 0.0001. The training procedure includes 100 epochs of pre-training followed by 200 epochs of fine-tuning. Balancing coefficients are empirically set as $\alpha = 0.1$ and $\beta = 0.01$ for all datasets. The model is implemented using the PyTorch framework and trained on an NVIDIA RTX 3080 GPU with 32 GB RAM.

Experimental Results and Analysis

Table 2 shows the performance of DIMVC-HIA across missing-view ratios $\eta \in [0.1, 0.7]$. The model outperforms

Dataset	Method	$\eta = 0.1$			$\eta = 0.3$			$\eta = 0.5$			$\eta = 0.7$		
		ACC	NMI	PUR	ACC	NMI	PUR	ACC	NMI	PUR	ACC	NMI	PUR
BDGP	COMPLETER	89.88	64.81	89.88	64.60	42.66	65.72	57.32	34.63	58.80	50.80	31.01	52.04
	CPSPAN	65.00	65.10	71.52	75.92	65.82	75.92	63.64	54.27	69.68	72.48	63.07	72.48
	DCG	95.38	87.00	95.38	91.59	76.89	91.59	76.47	58.03	76.47	66.00	47.23	66.00
	DIVIDE	89.68	64.61	89.68	64.40	42.46	65.52	57.12	34.43	58.60	51.60	30.81	51.84
	RPCIC	88.08	63.01	88.08	62.80	40.86	63.92	55.52	32.83	57.00	50.00	29.21	50.24
	APADC	65.52	55.89	65.52	82.48	68.28	82.48	67.16	56.14	67.60	64.36	51.10	64.36
	GIMVC	72.21	60.95	73.74	72.10	57.85	72.98	69.92	57.28	72.04	63.48	50.93	65.44
	PMIMC	87.96	72.14	87.96	89.52	75.55	89.52	93.04	80.92	93.04	<u>91.72</u>	77.78	<u>91.72</u>
	ProImp	75.20	70.21	77.12	92.56	82.20	92.56	90.92	77.76	90.92	84.72	65.63	84.72
	DSIMVC	<u>98.00</u>	<u>93.50</u>	<u>98.00</u>	<u>96.08</u>	<u>88.56</u>	<u>96.08</u>	<u>93.56</u>	<u>84.04</u>	<u>93.56</u>	91.12	<u>79.00</u>	91.12
	DIMVC-HIA	98.40	94.49	98.40	96.25	89.97	96.25	95.16	85.14	95.16	92.32	79.50	92.32
MNIST-USPS	COMPLETER	78.20	81.37	81.76	71.06	68.74	71.34	64.20	58.89	64.38	51.08	46.19	51.66
	CPSPAN	81.52	80.45	84.58	65.74	72.42	74.96	74.36	78.40	79.32	85.16	80.68	85.18
	DCG	<u>99.05</u>	<u>97.13</u>	<u>99.05</u>	<u>97.48</u>	<u>93.21</u>	<u>97.48</u>	96.09	90.11	96.09	92.58	83.17	92.58
	DIVIDE	<u>75.80</u>	<u>78.57</u>	<u>79.46</u>	<u>73.03</u>	69.54	<u>73.31</u>	63.90	58.59	64.08	50.78	45.89	51.36
	RPCIC	77.05	80.37	81.06	69.36	67.04	69.64	62.50	57.19	62.68	49.38	44.49	49.96
	APADC	97.32	93.58	97.32	96.24	90.87	96.24	94.16	86.61	94.16	91.32	82.51	91.32
	GIMVC	77.63	74.37	79.02	76.05	70.93	77.19	74.55	66.03	75.13	74.18	60.73	74.18
	PMIMC	77.34	76.27	79.74	72.48	72.15	74.76	76.14	75.63	79.88	78.94	74.40	79.12
	ProImp	99.02	<u>97.56</u>	99.02	97.46	93.82	97.46	96.32	<u>91.10</u>	96.32	93.42	86.17	93.42
	DSIMVC	98.00	94.98	98.00	97.08	93.92	97.08	<u>96.44</u>	91.01	96.44	<u>92.76</u>	<u>85.29</u>	<u>92.76</u>
	DIMVC-HIA	99.10	97.61	99.10	97.54	<u>93.84</u>	97.54	96.48	91.20	96.48	93.66	82.16	93.66
Fashion-MNIST	COMPLETER	70.55	84.92	78.91	64.92	68.27	68.09	60.64	61.36	61.80	55.41	49.24	49.31
	CPSPAN	66.63	72.51	71.97	76.64	77.66	80.80	60.82	70.44	67.44	59.75	69.68	66.00
	DCG	95.96	87.87	95.96	92.68	82.55	92.68	90.15	82.88	90.15	84.46	76.10	84.46
	DIVIDE	70.45	84.82	78.81	64.82	68.17	67.99	60.54	61.26	61.70	55.31	49.14	49.21
	RPCIC	68.85	83.22	77.21	63.22	66.57	66.39	58.94	59.66	60.10	43.71	47.54	47.61
	APADC	66.10	77.08	67.80	60.58	73.38	65.41	57.08	69.23	57.89	64.63	67.77	66.83
	GIMVC	75.98	77.87	78.98	75.47	74.46	78.36	69.65	69.09	73.63	67.46	64.86	71.37
	PMIMC	71.03	77.25	75.66	71.81	76.39	75.23	70.22	76.01	74.88	73.57	75.38	75.66
	ProImp	<u>96.26</u>	<u>92.45</u>	<u>96.26</u>	<u>93.48</u>	<u>87.89</u>	<u>93.48</u>	<u>91.01</u>	<u>83.99</u>	<u>91.01</u>	<u>86.74</u>	<u>78.19</u>	<u>86.74</u>
	DSIMVC	81.88	81.77	81.88	82.39	79.62	82.39	74.76	75.83	74.76	81.15	75.93	81.15
	DIMVC-HIA	98.84	97.12	98.84	97.16	93.72	97.16	96.51	92.25	96.51	95.27	92.05	95.27
HandWritten	COMPLETER	86.10	85.59	87.40	76.15	75.52	76.30	67.60	63.99	68.05	61.70	57.49	62.10
	CPSPAN	49.95	58.28	55.65	50.50	57.62	55.15	42.70	52.47	48.95	43.35	52.00	48.60
	DCG	86.74	84.30	86.74	82.07	77.28	82.08	81.13	76.71	81.13	79.08	72.45	79.08
	DIVIDE	85.90	85.39	87.20	75.95	75.32	76.10	67.40	63.79	67.85	61.50	57.29	61.90
	RPCIC	84.30	83.79	85.60	74.35	73.72	74.50	65.80	62.19	66.25	59.90	55.69	60.30
	APADC	80.75	83.02	82.95	82.35	83.29	83.95	83.35	<u>84.35</u>	83.90	78.05	75.49	79.75
	GIMVC	<u>92.14</u>	<u>88.34</u>	<u>93.02</u>	<u>93.58</u>	<u>87.03</u>	<u>93.58</u>	<u>90.73</u>	84.03	91.21	86.10	80.41	87.50
	PMIMC	81.55	78.88	81.80	91.75	85.21	91.75	78.65	78.79	82.20	77.65	78.26	80.70
	ProImp	82.25	78.79	82.25	80.10	76.80	80.10	81.00	77.56	81.00	78.90	73.29	78.90
	DSIMVC	75.55	71.30	79.15	73.10	70.53	75.70	71.00	67.87	75.70	65.35	64.44	69.00
	DIMVC-HIA	96.85	92.82	96.85	96.35	91.87	96.35	95.15	90.41	95.15	94.05	88.33	94.05

Table 2: Performance comparison with state-of-the-art methods across different datasets with varying missing rates η . The best results are bold, and the second-best results are underlined.

most of the ten baselines, with its advantage becoming more pronounced as incompleteness increases. At $\eta = 0.7$, it achieves 95.27% ACC on Fashion (8.53% higher than ProImp) and 94.05% on HandWritten (7.95% higher than GIMVC). On BDGP and MNIST-USPS, it maintains the best overall performance with an accuracy variation of only around 6%, demonstrating strong stability. In comparison, baseline methods such as DSIMVC experience significant drops in performance; for instance, its ACC on BDGP de-

creases from 98.00% to 91.12% as η increases. These results confirm the effectiveness of DIMVC-HIA’s hierarchical imputation in enabling robust and discriminative representation learning under severe view incompleteness.

Ablation Study

To further evaluate the effectiveness of the proposed DIMVC-HIA method, we performed an ablation study on the Fashion-MNIST dataset set with missing ratios (η) set

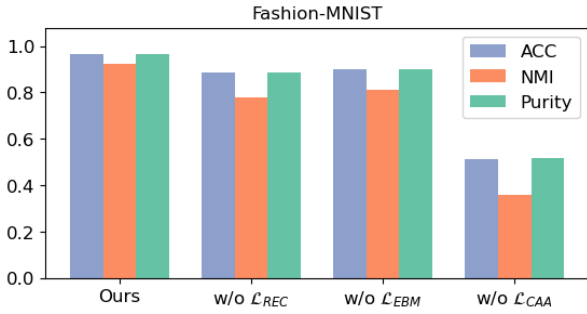


Figure 2: Ablation study results on the Fashion-MNIST dataset with $\eta = 0.5$.

to 0.5. We selectively removed key components including the reconstruction loss (\mathcal{L}_{REC}), the energy-based alignment loss (\mathcal{L}_{EBM}), and the contrastive assignment alignment loss (\mathcal{L}_{CAA}) in Eq. (23). The results are presented in Figure 2. Notably, removing the contrastive assignment alignment loss (\mathcal{L}_{CAA}) causes the most significant performance drop, highlighting its critical role in aligning clustering assignments. Meanwhile, excluding either the energy-based alignment loss (\mathcal{L}_{EBM}) or the reconstruction loss (\mathcal{L}_{REC}) also results in considerable degradation, confirming their importance in ensuring stable optimization and learning robust, discriminative representations.

Visualization Results and Convergence Analysis

To evaluate the clustering performance of DIMVC-HIA, Fig. 3 presents the T-SNE visualization of the complete feature embeddings (H^*) on the BDGP and Handwritten datasets, with missing ratios of $\eta = 0.5$ and $\eta = 0.7$, respectively. The t-SNE visualization on the BDGP dataset reveals distinct and compact clusters with well-defined inter-class boundaries, aligning with its high clustering accuracy. The clear separation and minimal overlap between clusters reflect the model’s strong discriminative capability. Despite the high missing ratio and a larger number of clusters, the t-SNE result on the Handwritten dataset also exhibits visibly separable cluster structures.

To evaluate the convergence behavior of DIMVC-HIA, Figure 4 shows the loss curves on the MNIST-USPS and Fashion-MNIST datasets under $\eta = 0.1$. The loss exhibits a consistent downward trend throughout training, reflecting stable convergence. In both datasets, a rapid decline occurs within the first 25 epochs, followed by a slower, steady reduction that eventually converges to a stable value. This pattern demonstrates the effectiveness of DIMVC-HIA’s optimization strategy and highlights the robustness of its architecture in facilitating efficient and reliable training.

Parameter Sensitivity Analysis

In DIMVC-HIA, two trade-off hyperparameters, α and β , are introduced to balance multiple objectives in the loss function. To assess the model’s sensitivity to these parameters, we conducted experiments on the BDGP and Handwritten datasets under a missing ratio (η) of 0.3. As




Figure 3: T-SNE visualizations of embeddings on BDGP with $\eta = 0.5$ (left) and Handwritten with $\eta = 0.7$ (right).

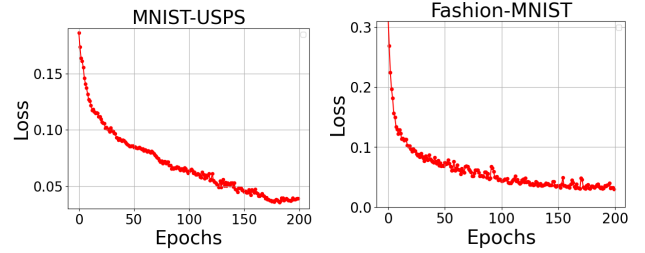


Figure 4: Convergence curves for the MNIST-USPS and Fashion-MNIST datasets with $\eta = 0.1$.

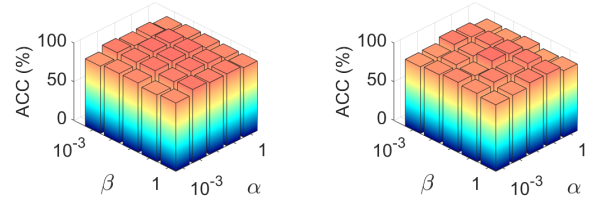


Figure 5: Parameter sensitivity of α and β on BDGP (left) and Handwritten (right) datasets with $\eta = 0.3$, where both α and β take values in $\{0.001, 0.01, 0.05, 0.1, 1\}$.

shown in Figure 5, the clustering performance (measured by ACC) remains consistently strong as α varies in the range of $[0.01, 0.10]$ and β in $[0.01, 0.05]$. The results indicate that the overall performance is relatively stable, exhibiting no substantial fluctuations across the tested configurations. Consequently, we choose $\alpha = 0.1$ and $\beta = 0.01$ as the default settings for all datasets.

Conclusion

This paper presents DIMVC-HIA, a hierarchical imputation and alignment framework for incomplete multi-view clustering, which integrates view-specific autoencoders, energy-based alignment, and contrastive cluster assignment to ensure reliable imputation, semantic consistency, and compact clustering under varying levels of view incompleteness. Extensive experiments demonstrate DIMVC-HIA’s superiority over IMVC baseline methods across a range of missing ratios, with results from multiple evaluation metrics confirming its ability to produce well-structured and discriminative representations under severe view incompleteness.

Acknowledgements

This work was supported in part by the NSF (CNS-2153358, CNS-2245918), the DoD CoE-AIML (W911NF-20-2-0277), and the Commonwealth Cyber Initiative (CCI and COVA CCI).

References

- Asuncion, A.; Newman, D.; et al. 2007. UCI machine learning repository.
- Bachtis, D.; Biroli, G.; Decelle, A.; and Seoane, B. 2024. Cascade of phase transitions in the training of energy-based models. *Advances in neural information processing systems*, 37: 55591–55619.
- Bai, S.; Zheng, Q.; Ren, X.; and Zhu, J. 2024. Graph-guided imputation-free incomplete multi-view clustering. *Expert Systems with Applications*, 258: 125165.
- Cai, X.; Wang, H.; Huang, H.; and Ding, C. 2012. Joint stage recognition and anatomical annotation of drosophila gene expression patterns. *Bioinformatics*, 28(12): i16–i24.
- Cao, J.; Hu, Y.; Tan, Z.; and Zhao, X. 2025. Cross-modal Multi-task Learning for Multimedia Event Extraction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 11454–11462.
- Chao, G.; Jiang, Y.; and Chu, D. 2024. Incomplete contrastive multi-view clustering with high-confidence guiding. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 11221–11229.
- Chen, H.; Xu, C.; Guan, Z.; Zhao, W.; and Liu, J. 2025. Biased Incomplete Multi-View Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 15767–15775.
- Dai, Y.; Jin, J.; Dong, Z.; Wang, S.; Liu, X.; Zhu, E.; Yang, X.; Gan, X.; and Feng, Y. 2025. Imputation-free and Alignment-free: Incomplete Multi-view Clustering Driven by Consensus Semantic Learning. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 5071–5081.
- Ding, X.; and Yang, F. 2024. Multi-view randomized kernel classification via nonconvex optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 11793–11801.
- Dong, Z.; Hu, D.; Jin, J.; Wang, S.; Liu, X.; and Zhu, E. 2025. Selective Cross-view Topology for Deep Incomplete Multi-view Clustering. *IEEE Transactions on Image Processing*.
- Dou, F.; Lu, J.; Zhu, T.; and Bi, J. 2025. LOCAL: Latent Orthonormal Contrastive Learning for Paired Image Classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6408–6417.
- Du, Y.; Wang, Y.; Wang, Z.; Ning, R.; and Li, L. 2025. PGFormer: A Prototype-Graph Transformer for Incomplete Multiview Clustering. *IEEE Transactions on Neural Networks and Learning Systems*.
- Gu, Z.; Feng, S.; Li, Z.; Yuan, J.; and Liu, J. 2024. NOODLE: Joint cross-view discrepancy discovery and high-order correlation detection for multi-view subspace clustering. *ACM Transactions on Knowledge Discovery from Data*, 18(6): 1–23.
- Guan, R.; Tu, W.; Wang, S.; Liu, J.; Hu, D.; Tang, C.; Feng, Y.; Li, J.; Xiao, B.; and Liu, X. 2025. Structure-adaptive multi-view graph clustering for remote sensing data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 16933–16941.
- Ji, J.; and Feng, S. 2025. Anchors Crash Tensor: Efficient and Scalable Tensorial Multi-view Subspace Clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11600–11609.
- Kim, J.; Lee, K.; and Park, T. 2025. To Predict or Not to Predict? Proportionally Masked Autoencoders for Tabular Data Imputation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 17886–17894.
- LeCun, Y.; Chopra, S.; Hadsell, R.; Ranzato, M.; Huang, F.; et al. 2006. A tutorial on energy-based learning. *Predicting structured data*, 1(0).
- Li, C.; Xu, C.; and Lu, Y. 2019. Incomplete Multi-view Clustering via GAN. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 11183–11192.
- 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*, 3911–3919.
- Li, Y.; Zhao, T.; and Qin, J. 2021. Contrastive Multi-View Representation Learning for Incomplete Data. *Neural Networks*, 144: 191–203.
- Lin, Y.; Gou, Y.; Liu, Z.; Li, B.; Lv, J.; and Peng, X. 2021. Completer: Incomplete multi-view clustering via contrastive prediction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11174–11183.
- Liu, C.; Jia, J.; Wen, J.; Liu, Y.; Luo, X.; Huang, C.; and Xu, Y. 2024a. Attention-induced embedding imputation for incomplete multi-view partial multi-label classification. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 13864–13872.
- Liu, J.; Liu, X.; Yang, Y.; Liao, Q.; and Xia, Y. 2023. Contrastive multi-view kernel learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(8): 9552–9566.
- Liu, S.; Zhang, J.; Wen, Y.; Yang, X.; Wang, S.; Zhang, Y.; Zhu, E.; Tang, C.; Zhao, L.; and Liu, X. 2024b. Sample-level cross-view similarity learning for incomplete multi-view clustering. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 14017–14025.

- 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 *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 14193–14201.
- Margeloiu, A.; Jiang, X.; Simidjievski, N.; and Jamnik, M. 2024. Tabebm: A tabular data augmentation method with distinct class-specific energy-based models. *Advances in Neural Information Processing Systems*, 37: 72094–72144.
- Nie, W.; Vahdat, A.; and Anandkumar, A. 2021. Controllable and compositional generation with latent-space energy-based models. *Advances in Neural Information Processing Systems*, 34: 13497–13510.
- Park, K.; and Lee, S. 2025. SMMF: Square-Matricized Momentum Factorization for Memory-Efficient Optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 19848–19856.
- Peng, T.; Li, Z.; Wang, P.; Zhang, L.; and Zhao, H. 2024. A novel energy based model mechanism for multi-modal aspect-based sentiment analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 18869–18878.
- 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 *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 14633–14641.
- Salakhutdinov, R.; and Hinton, G. 2009. Deep boltzmann machines. In *Artificial intelligence and statistics*, 448–455. PMLR.
- Shiri, M.; Reddy, M. P.; and Sun, J. 2024. Supervised contrastive vision transformer for breast histopathological image classification. In *2024 IEEE International Conference on Information Reuse and Integration for Data Science (IRI)*, 296–301. IEEE.
- Sun, L.; Huang, Z.; Wang, Z.; Wang, F.; Peng, H.; and Yu, P. S. 2024. Motif-aware riemannian graph neural network with generative-contrastive learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 9044–9052.
- Tan, Y.; Cai, H.; Huang, S.; Wei, S.; Yang, F.; and Lv, J. 2024. An effective augmented lagrangian method for fine-grained multi-view optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 15258–15266.
- Tang, H.; and Liu, Y. 2022. Deep safe incomplete multi-view clustering: Theorem and algorithm. In *International conference on machine learning*, 21090–21110. PMLR.
- Tu, W.; Guan, R.; Zhou, S.; Ma, C.; Peng, X.; Cai, Z.; Liu, Z.; Cheng, J.; and Liu, X. 2024. Attribute-missing graph clustering network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 15392–15401.
- Wang, X.; Li, Z.; Liu, Y.; and Wu, Q. 2025a. Deep Incomplete Multi-view Clustering via Contrastive Alignment and Imputation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(2): 1234–1248.
- Wang, Z.; Du, Y.; Ning, R.; and Li, L. 2025b. Energy-based Deep Incomplete Multi-View Clustering. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 1686–1694.
- Wang, Z.; Du, Y.; Wang, Y.; Ning, R.; and Li, L. 2025c. Deep incomplete multi-view clustering via multi-level imputation and contrastive alignment. *Neural Networks*, 181: 106851.
- Xiao, H.; Rasul, K.; and Vollgraf, R. 2017. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*.
- Xu, G.; Wen, J.; Liu, C.; Hu, B.; Liu, Y.; Fei, L.; and Wang, W. 2024. Deep variational incomplete multi-view clustering: Exploring shared clustering structures. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 16147–16155.
- Xu, J.; Li, C.; Peng, L.; Ren, Y.; Shi, X.; Shen, H. T.; and Zhu, X. 2023. Adaptive feature projection with distribution alignment for deep incomplete multi-view clustering. *IEEE Transactions on Image Processing*, 32: 1354–1366.
- Xue, Y.; Zhao, Y.; Liu, J.; and Yin, J. 2022. Robust Incomplete Multi-view Clustering with Diversified Graph Contrastive Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 1534–1542.
- Yang, X.; Sun, Y.; Chen, X.; Zhang, Y.; and Yuan, X. 2025. Graph Structure Learning for Spatial-Temporal Imputation: Adapting to Node and Feature Scales. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 959–967.
- Yuan, H.; Lai, S.; Li, X.; Dai, J.; Sun, Y.; and Ren, Z. 2024. Robust Prototype Completion for Incomplete Multi-view Clustering. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 10402–10411.
- 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 Transactions on Image Processing*.
- Zang, Y.; Ren, L.; Li, Y.; Wang, Z.; Selby, D. A.; Wang, Z.; Vollmer, S. J.; Yin, H.; Song, J.; and Wu, J. 2025. Rethinking cancer gene identification through graph anomaly analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 13161–13169.
- Zeng, X.; Peng, H.; and Li, A. 2024. Adversarial socialbots modeling based on structural information principles. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 392–400.
- Zhang, L.; Wang, Q.; and Yang, T. 2023. A Unified Prototype-based Framework for Incomplete Multi-view Clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12415–12424.
- Zhang, Y.; Lin, Y.; Yan, W.; Yao, L.; Wan, X.; Li, G.; Zhang, C.; Ke, G.; and Xu, J. 2025. Incomplete Multi-view Clustering via Diffusion Contrastive Generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 22650–22658.