TIGER: TARGET-AWARE INTERVENTIONAL GRAPH-STRUCTURE DECOUPLING AND RETRIEVAL FOR INCOMPLETE MULTI-VIEW CLUSTERING

Anonymous authors
Paper under double-blind review

ABSTRACT

Incomplete multi-view clustering (IMVC) aims to learn effective clustering representations from multi-view data with partially missing views. Although considerable progress has been made, existing methods still face several challenges: 1) Many approaches recover missing views through direct feature reconstruction or global information sharing, which may ignore the distribution characteristics of the target view and lead to inaccurate representation generation. 2) The structural relationships among samples and the complementary information across views are insufficiently explored during missing-view recovery, making the learned representations sensitive to incomplete and noisy observations. 3) Existing methods usually lack an explicit mechanism to evaluate the reliability of recovered information, causing uncertain completion results to negatively affect representation fusion and clustering optimization. To address these issues, we propose TIGER: Target-aware Interventional Graph-structure Decoupling and Retrieval for Incomplete Multi-view Clustering. Specifically, TIGER first learns view-specific latent representations and separates transferable and view-dependent information to construct more reliable cross-view relationships. Then, a target-aware retrieval strategy is developed to recover missing views by selecting informative anchors from observed target-view samples according to cross-view structural evidence. Furthermore, an intervention-based completion mechanism is introduced to identify and reduce unreliable local correlations, enabling robust representation recovery under severe missing conditions. Finally, the recovered representations are integrated with clustering learning through a unified optimization framework. Extensive experiments on seven benchmark datasets under various missing rates demonstrate that TIGER consistently outperforms state-of-the-art IMVC methods.

1 INTRODUCTION

Multi-view clustering (MVC) discovers latent semantic groups by integrating complementary observations from multiple feature spaces or modalities. Recent deep MVC methods increasingly exploit cross-view correspondence, neighborhood relations, and semantic–structural cues rather than relying on point-wise feature agreement Eisenberg et al. (2025); Zhou et al. (2026); Peng et al. (2026); Kong et al. (2026). In practice, sensor failures, privacy restrictions, communication loss, and acquisition costs often make entire views unavailable, requiring incomplete multi-view clustering (IMVC) to preserve cross-view semantics while remaining compatible with missing target-view geometry. Structure-oriented IMVC methods show that reliable neighborhoods and relational evidence can stabilize learning under missingness Wen et al. (2023); Yu et al. (2024), while diffusion-based recovery demonstrates effective generative completion at high missing rates Wen et al. (2024).

Recent studies couple missing-view completion with clustering-oriented representation learning. Scalable anchor learning and adaptive similarity reconstruction improve structural recovery without dense pairwise graphs Qin et al. (2025); Xu et al. (2025), while deep imputation and hierarchical semantic alignment transfer information across incomplete views Liu et al. (2025); Ding et al. (2025). Federated graph guidance extends structural transfer to distributed settings Chao et al. (2025), while cyclic variational autoencoders model cross-view correspondence Gao & Pu (2025). In contrast, FreeCSL avoids explicit imputation and learns consensus semantics from observations Dai et al.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

(2025). Despite these advances, most recovery relies on source-side representations, global prototypes, or shared latent spaces, leaving observed target-view geometry only indirectly exploited.

More importantly, existing structural guidance is transferred as an entangled relational signal, where neighborhood graphs may mix transferable semantics with view-dependent variations caused by modality-specific distributions, noise, or local distortions. Directly using such relations for completion can therefore propagate unreliable view-specific dependencies together with useful structural consensus. Recent methods emphasize geometric consistency and reliable information transfer: GAVIM preserves shared information and local geometry in an imputation-free variational framework Chen et al. (2026b), while IMVC-NCD conditions diffusion-based recovery on neighborhood structure Guo et al. (2026); federated tensor constraints and prototype–sample alignment further improve global consistency Feng et al. (2026); Zhu et al. (2026). Nevertheless, they do not explicitly disentangle transferable graph topology from view-specific relational variation before recovery, allowing unreliable dependencies to propagate across views. Reliable IMVC requires determining which structural relations should be transferred, rather than merely exploiting graph structure.

To address the limitations, we propose **TIGER**: **T**arget-aware **I**nterventional **G**raph-structure d**E**coupling and **R**etrieval for Incomplete Multi-view Clustering. TIGER factorizes representations into stable and private components, decoupling graph structure by aligning transferable topology while separating view-dependent variations. For missing target views, stable evidence retrieves target-bank anchors, grounding recovery in target geometry over source-side features. Retrieval confidence controls residual correction, while graph interventions challenge fragile private dependencies before reliability-aware fusion and clustering. Our contributions are summarized as follows:

- We propose a target-aware graph retrieval framework for IMVC using transferable structural evidence to locate anchors in the target-view bank, reformulating missing-view recovery as target-grounded retrieval rather than source-to-target feature reconstruction.
- To the best of our knowledge, we are the first to introduce graph-structure decoupling into incomplete multi-view clustering. Rather than treating structure as an indivisible signal, TIGER aligns transferable stable topology across views while separating view-specific private variation, providing reliable structural guidance for missing-view recovery.
- We introduce confidence-guided stochastic graph intervention and residual completion, which challenge fragile private relations and adaptively control correction based on retrieval reliability, reducing the impact of uncertain completion on fusion and clustering.

2 Related Works

Deep Incomplete Multi-View Clustering. Deep incomplete multi-view clustering (DIMVC) learns representations from partially observed views while reducing distortion caused by missingness. CPSPAN aligns partial samples and prototypes across views (Jin et al., 2023), whereas ICMVC transfers consistency relations with high-confidence guidance (Chao et al., 2024). Methods couple recovery with latent modeling: DVIMC-ITG introduces information-theoretic variational learning (Chen et al., 2025), DDA-IMVC performs dual-level distribution alignment (Ren et al., 2025), and GAVIM preserves information and view geometry without explicit imputation (Chen et al., 2026b). DIMVC-HIA combines hierarchical imputation with semantic alignment (Du et al., 2026), while dynamic graph learning reconstructs masked structures during clustering (Zhang et al., 2026b). AFIMVC, RAC-DMVC, and EFDMVC improve robustness through adversarial disentanglement, reliability-aware contrastive learning, and uncertainty-aware federated fusion (Wang et al., 2026b; Dong et al., 2026; Wei et al., 2026), while IDCD disentangles semantic and view-specific factors with conditional diffusion (Chen et al., 2026c). Despite these advances, existing methods remain weakly grounded in target-view geometry, whereas our method retrieves target-view anchors from transferable structure and challenges unreliable private relations under severe incompleteness.

Deep Multi-View Feature Learning. Deep multi-view feature learning seeks representations preserving cross-view consensus and view-specific diversity across heterogeneous views and semantic distributions. DIB learns deterministic features via an information bottleneck (Yan et al., 2024), while SURER propagates structural relations through a graph neural network (Wang et al., 2024). DMVC-CE uses collaborative experts for representation learning (Zhang et al., 2025c), DMAC learns adaptive anchors via graph convolution and mutual-information

guidance (Wang et al., 2025), and AIRMVC rectifies noisy observations before contrastive learning (Yang et al., 2025). Methods emphasize structured interaction: DC-SPAN decomposes shared and private subspaces (Chen et al., 2026a), DCL-MVC applies contrastive objectives at fusion and representation levels (Liu et al., 2026), MGLC models global-local dependencies through sequence learning (Zhang et al., 2026a), and adaptive evolutionary fusion searches hierarchical fusion structures (Zhao et al., 2026). However, these advances primarily improve representation interaction rather than recovering target-view structure under missing observations. They strengthen consensus-diversity modeling but assume available view features and focus on alignment or fusion, whereas our method converts decoupled structure into retrieval evidence for missing target representations.

3 METHOD

3.1 PROBLEM DEFINITION

Given an incomplete multi-view dataset with V views and N samples, view v is denoted by $X^v \in \mathbb{R}^{N\times d^v}$, where d^v is its feature dimensionality. An observation indicator matrix $M = [M_i^v] \in \{0,1\}^{N\times V}$ describes data availability, where $M_i^v = 1$ indicates sample i is observed in view v, and $M_i^v = 0$ indicates the view is missing. Given the number of clusters K, incomplete multi-view clustering aims to learn a unified representation that preserves consensus semantics while retaining view-specific information from partially observed data. The framework is illustrated in Figure 1.

Definition 3.1. *For sample i, the observed-view set and missing-view set are $\Omega_i = \{v \mid M_i^v = 1\}$ and $\bar{\Omega}_i = \{v \mid M_i^v = 0\}$, respectively. For view v, the observed-sample index set is denoted by $\mathcal{I}_v = \{i \mid M_i^v = 1\}$. We assume that every sample retains at least one observed view, i.e., $|\Omega_i| \geq 1$.*

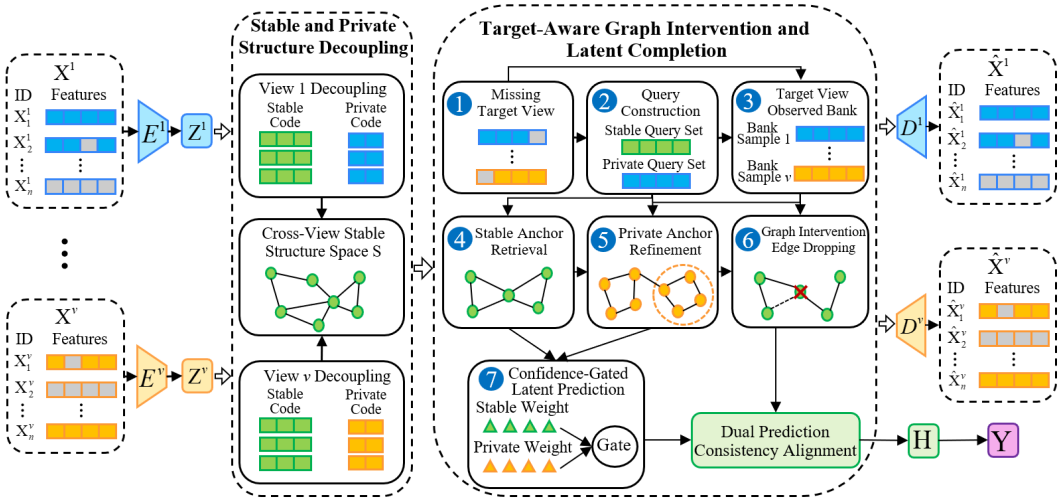

Figure 1: Overall framework of TIGER. Observed views are encoded into latent representations, yielding stable and private factors for cross-view structural relations. For a missing target view, TIGER aggregates stable evidence into a target-aware query, retrieves stable and private anchors from the target-view bank, perturbs unreliable private edges via stochastic graph intervention, and controls residual correction by retrieval confidence. Reliable representations are fused for clustering.

3.2 VIEW REPRESENTATION AND GRAPH-STRUCTURE DECOUPLING

Different views usually exhibit heterogeneous scales, distributions, and noise patterns. Directly establishing cross-view relations in the original feature space can therefore mix semantic agreement with view-specific disturbances. To provide a more reliable and stable basis for subsequent graph modeling, we first extract a compact latent representation within each observed view.

For each observed sample, the encoder and decoder mappings are written as:

$$z_i^v = E^v(x_i^v), \qquad \hat{x}_i^v = D^v(z_i^v), \qquad i \in \mathcal{I}_v, \tag{1}$$

where $E^v(\cdot)$ and $D^v(\cdot)$ denote the encoder and decoder of view v, respectively; z_i^v is the latent representation of sample i in view v; and $\hat{x}_i^v$ is the corresponding reconstruction.

To prevent the latent space from being dominated by a small number of large reconstruction errors, reconstruction is constrained only on observed samples using a smooth L_1 penalty:

$$L_{\text{rec}} = \frac{1}{\sum_{v=1}^{V} |\mathcal{I}_v|} \sum_{v=1}^{V} \sum_{i \in \mathcal{I}_v} \ell_h(\hat{x}_i^v - x_i^v), \tag{2}$$

where $\ell_h(\cdot)$ is the smooth L_1 loss averaged over feature dimensions, e denotes the reconstruction residual, and $\delta > 0$ controls the transition between quadratic and linear regions. Specifically, $\ell_h(e) = e^2/(2\delta)$ for $|e| \leq \delta$ and $|e| - \delta/2$ otherwise. We set $\delta = 1.0$ throughout experiments, yielding smooth optimization for small residuals while limiting the influence of large errors.

A single latent code still mixes cross-view consensus with view-specific variations. We therefore factorize each latent representation into stable and private components for better separation:

$$s_i^v = F_s^v(z_i^v), \qquad p_i^v = F_p^v(z_i^v), \tag{3}$$

where $F_s^v(\cdot)$ and $F_p^v(\cdot)$ denote the stable and private branches, respectively; s_i^v captures information expected to be transferable across views, while p_i^v retains complementary view-specific variations.

Recent deep multi-view clustering work has shown that explicitly separating shared and view-specific pathways can alleviate representation entanglement (Xiao et al., 2026). Inspired by this principle, TIGER uses stable-private factorization only as a representation basis and places its main emphasis on further decoupling the relational graph structure for robust target-view recovery under severe incompleteness rather than treating feature factorization itself as the key contribution.

Specifically, stable information should preserve similar local topology across views, whereas private information should not encode the same stable directions. Instead of enforcing point-wise feature equality, we convert stable representations into neighborhood graphs. For view pair (v, u), let $\mathcal{I}_{vu} = \mathcal{I}_v \cap \mathcal{I}_u$ denote the jointly observed set. The local neighborhood distribution of sample i in view v is defined as:

$$\pi_{ij}^{vu} = \frac{\exp\big(\cos(s_i^v, s_j^v)/\tau\big)}{\sum_{k \in \mathcal{I}_{vu},\, k \neq i} \exp(\cos(s_i^v, s_k^v)/\tau)}, \qquad j \in \mathcal{I}_{vu},\ j \neq i, \tag{4}$$

where τ is a temperature parameter and $\cos(\cdot,\cdot)$ denotes cosine similarity. Each vector $\boldsymbol{\pi}_i^{vu}$ can be viewed as a row of a sample-wise soft relation graph, so comparing these distributions directly measures whether two views preserve the same local topology around sample i.

Based on these relation graphs, we define the graph-structure decoupling objective as:

$$L_{\text{str}} = \frac{1}{|\mathcal{P}|} \sum_{(v,u) \in \mathcal{P}} \frac{1}{|\mathcal{I}_{vu}|} \sum_{i \in \mathcal{I}_{vu}} \text{JSD}(\boldsymbol{\pi}_i^{vu}, \boldsymbol{\pi}_i^{uv}) + \mu \frac{1}{\sum_{v=1}^{V} |\mathcal{I}_v|} \sum_{v=1}^{V} \sum_{i \in \mathcal{I}_v} |\cos(s_i^v, p_i^v)|, \tag{5}$$

where $\mathcal{P} = \{(v,u) \mid 1 \leq v < u \leq V, |\mathcal{I}_{vu}| > 0\}$ is the set of valid view pairs; $\boldsymbol{\pi}_i^{vu}$ is the neighborhood-distribution vector composed of $\{\pi_{ij}^{vu}\}$; $\text{JSD}(\cdot,\cdot)$ denotes the Jensen–Shannon divergence; and μ controls the stable-private decorrelation strength. The first term aligns cross-view stable topology at the neighborhood level, while the second prevents private factors from collapsing back onto the stable subspace. In this way, TIGER separates transferable graph relations from view-specific variation before missing-view recovery in a noise-aware and robust manner.

3.3 Target-Aware Cross-View Anchor Retrieval

After graph-structure decoupling, the remaining challenge is how to transfer reliable information into a missing target view without directly copying the bias of a particular source view. Recent graph-based IMVC research shows that explicitly exploiting cross-view compensation relationships can improve the recovery of missing subgraph structures in a targeted way (Wang et al., 2026c). Motivated by this structural compensation perspective, TIGER constructs a sample-specific target retrieval graph: stable evidence from all available source views determines where to search, while the observed samples in the target view provide the actual anchors for recovery.

For a missing target view $v \in \bar{\Omega}_i$ of sample i, stable structures from all observed views form a normalized target-aware query:

$$q_i^v = \frac{\sum_{u \in \Omega_i} s_i^u}{\left\| \sum_{u \in \Omega_i} s_i^u \right\|_2}, \tag{6}$$

where q_i^v is the stable query for target view v, and $\| \cdot \|_2$ denotes the L_2 norm. Aggregating all available stable factors reduces dependence on any single source view and provides a view-agnostic structural cue for reliably locating the target manifold under missing-view conditions.

The query is then compared with stable structures in the observed target-view bank under the target-view latent geometry, and the top k samples are selected to form the local candidate set:

$$\mathcal{A}_i^v = \text{Top-k}_{j \in \mathcal{I}_v} \cos(q_i^v, s_j^v), \tag{7}$$

where $\mathcal{A}_i^v$ denotes the candidate set of sample i in target view v, and k the neighborhood size. We set $k = 36$ throughout experiments. This step converts cross-view evidence into a compact target-specific retrieval graph, restricting completion to informative target samples and excluding potentially unrelated observations, preserving reliable local structural evidence for recovery.

To determine the main position of the missing sample on the target manifold, stable retrieval weights are computed within the candidate set and used to aggregate the corresponding target latent codes:

$$\alpha_{ij}^v = \frac{\exp\big(\cos(q_i^v, s_j^v)/\tau_r\big)}{\sum_{k \in \mathcal{A}_i^v} \exp(\cos(q_i^v, s_k^v)/\tau_r)}, \qquad r_i^v = \sum_{j \in \mathcal{A}_i^v} \alpha_{ij}^v z_j^v, \tag{8}$$

where α_{ij}^v measures the stable contribution of candidate j to the target query, τ_r is the stable retrieval temperature, and r_i^v is the stable anchor. The stable anchor provides a consensus-dominated estimate of the target location and serves as the reliable prediction baseline in the next stage.

Stable structure alone, however, may smooth away fine-grained characteristics that are specific to the target view. We therefore refine the retrieval graph with private information from the same target candidates:

$$\beta_{ij}^v = \frac{\exp\big(\cos(F_p^v(r_i^v), p_j^v)/\tau_p\big)}{\sum_{k \in \mathcal{A}_i^v} \exp\big(\cos(F_p^v(r_i^v), p_k^v)/\tau_p\big)}, \qquad u_i^v = \sum_{j \in \mathcal{A}_i^v} \beta_{ij}^v p_j^v, \tag{9}$$

where β_{ij}^v is the private matching weight, τ_p is the private retrieval temperature, and u_i^v is the private anchor. The stable and private anchors therefore play complementary roles during completion: the former fixes the principal target position, while the latter supplies target-specific local variation.

3.4 Confidence-Guided Graph Intervention and Latent Completion

Not every retrieval graph is equally reliable. When the stable weights are scattered across many candidates, the target position is ambiguous and a strong private correction can amplify retrieval errors. High-confidence guidance has been shown to improve the reliability of incomplete multi-view graph clustering (Zhao et al., 2025). Inspired by this reliability-aware principle, we construct a robust sample-level confidence directly from the concentration of stable retrieval weights, rather than relying on clustering pseudo-labels. The retrieval confidence is then defined as:

$$c_i^v = 1 + \frac{\sum_{j \in \mathcal{A}_i^v} \alpha_{ij}^v \log \alpha_{ij}^v}{\log k}, \tag{10}$$

where $c_i^v \in [0, 1]$ is the retrieval confidence for sample i in target view v. The numerator is the negative entropy of the stable retrieval distribution. When one or a few candidates dominate, the entropy is low and c_i^v approaches 1; when the weights are nearly uniform, the entropy approaches its maximum $\log k$ and c_i^v becomes small. Hence, this confidence measures how decisively the cross-view stable evidence more precisely locates the missing sample in the target bank.

Confidence controls the magnitude of the later private correction, but it does not by itself remove accidental private edges. To explicitly test whether the completion relies on fragile local relations under repeated local perturbations, TIGER keeps stable retrieval weights fixed and applies two independent stochastic interventions only to the private candidate graph. For the t-th intervention:

$$\bar{\beta}_{ij}^{(v,t)} = \frac{e_{ij}^{(v,t)} \beta_{ij}^v}{\sum_{k \in \mathcal{A}_i^v} e_{ik}^{(v,t)} \beta_{ik}^v}, \qquad u_i^{(v,t)} = \sum_{j \in \mathcal{A}_i^v} \bar{\beta}_{ij}^{(v,t)} p_j^v, \tag{11}$$

where $t \in \{1, 2\}$ is the intervention index, $e_{ij}^{(v,t)} \in \{0, 1\}$ indicates whether a candidate edge is retained in the t-th intervention, $\bar{\beta}_{ij}^{(v,t)}$ is the intervened private weight, and $u_i^{(v,t)}$ is the corresponding private anchor. At least one candidate edge is retained in each intervention. Because only private edges are perturbed, the stable anchor remains the fixed structural reference while unreliable target-specific relations are challenged, making the prediction less sensitive to local perturbations.

Using the stable anchor as a reliable prediction baseline, the model learns only a confidence-gated private residual term:

$$y_i^{(v,t)} = r_i^v + c_i^v G^v\left([r_i^v \,\|\, u_i^{(v,t)}]\right), \tag{12}$$

where $y_i^{(v,t)}$ is the target latent prediction under the t-th graph intervention, $G^v(\cdot)$ is the completion network, and $[\cdot\|\cdot]$ denotes vector concatenation. This design makes the prediction close to the stable anchor when retrieval is uncertain, while allowing stronger target-specific refinement in a sample-adaptive and reliability-aware manner when the retrieved neighborhood is reliable.

To train the completion module without labels, we mask samples that are originally observed in view v and have at least one additional observed view to construct self-supervised target-view training pairs, and denote their indices by $\mathcal{T}^v$. The graph-intervention completion objective is defined as:

$$L_{\text{gii}} = \frac{1}{2\sum_{v=1}^{V}|\mathcal{T}^v|}\sum_{v=1}^{V}\sum_{i\in\mathcal{T}^v}\sum_{t=1}^{2}\left\|y_i^{(v,t)} - z_i^v\right\|_1 + \eta\frac{1}{\sum_{v=1}^{V}|\mathcal{T}^v|}\sum_{v=1}^{V}\sum_{i\in\mathcal{T}^v}\left\|y_i^{(v,1)} - y_i^{(v,2)}\right\|_1, \tag{13}$$

where $\mathcal{T}^v$ is the self-supervised training set for view v, η weights the intervention-consistency term, and $\|\cdot\|_1$ denotes the L_1 norm, i.e., the sum of absolute errors over all dimensions. The first term teaches the cross-view retrieval process to recover the true target latent position, whereas the second suppresses prediction changes caused by different private-edge interventions. Their combination encourages both recovery accuracy and structural robustness under missing views.

3.5 Unified Representation and Overall Objective

After target-aware completion, the mean of the two intervention predictions is used as the completed latent representation for a missing view, whereas an observed view retains its original latent code. The stable branch is applied to all observed or completed views. For sample i, we denote the unified representation by h_i, and collect all sample representations as $H = \{h_i\}_{i=1}^N$. The view-wise stable representations are fused with reliability-aware weights: an observed view receives weight 1, while a completed view receives the retrieval confidence c_i^v. This prevents low-confidence completions from dominating H while allowing reliable recovered views to contribute complementary structure. Finally, the cluster assignment matrix Y is obtained by applying K-means to H.

Algorithm 1 Learning Procedure of TIGER

Require: Incomplete multi-view data $\mathcal{X} = \{X^v\}_{v=1}^V$, mask M, cluster number K.
Ensure: Clustering labels Y.

```
Initialize TIGER parameters and observation sets from M.
for joint training stage do
    Obtain z_i^v and x̂_i^v by Eq. (1), and compute L_rec by Eq. (2).
    Generate stable and private factors by Eq. (3).
    Construct cross-view neighborhood distributions by Eq. (4) and compute L_str by Eq. (5).
    Construct self-supervised pseudo-missing sets T^v.
    Form q_i^v and A_i^v by Eqs. (6)–(7).
    Retrieve r_i^v and u_i^v by Eqs. (8)–(9).
    Compute c_i^v by Eq. (10) and intervene on private edges by Eq. (11).
    Generate y_i^(v,t) by Eq. (12) and compute L_gii by Eq. (13).
    Minimize L in Eq. (14) to update all learnable parameters.
end for
Complete missing views using Eqs. (6)–(12) and average the two intervention predictions.
Fuse observed and completed views into H with reliability-aware weights.
Apply K-means to H to obtain Y.
return Y.
```

The three top-level objectives constrain complementary aspects of the model: reconstruction preserves observed-view content, graph-structure decoupling establishes transferable relational topology while suppressing stable-private redundancy, and graph-intervention completion learns robust target-view recovery under incomplete and noisy view conditions. Their joint objective is:

$$L = L_{\text{rec}} + \lambda_{\text{str}} L_{\text{str}} + \lambda_{\text{gii}} L_{\text{gii}}, \tag{14}$$

where λ_{str} and λ_{gii} balance the graph-structure decoupling and graph-intervention completion terms, respectively. Joint optimization in a unified manner allows the stable topology learned from observed data to guide target retrieval, while self-supervised completion improves the latent space for discriminative structural modeling through mutual refinement. TIGER's complete learning and inference procedure is summarized in Algorithm 1, with computational complexity analyzed in Appendix B.4.

4 Experiments

4.1 Experimental Setup

1) Datasets: We evaluate TIGER on seven benchmark datasets in incomplete multi-view clustering, including **Synthetic3D** (Dong et al., 2025), **HW** (Zhang et al., 2025a), **MNIST4** (Zhang et al., 2025b), **MNIST-Fea** (Zhang et al., 2025b), **CUB** (Gao & Pu, 2025), **WebKB** (Qin et al., 2025), and **ALOI-100** (Wang et al., 2026a). Following the standard setting, we randomly remove views at the designated missing rates while ensuring that every sample retains at least one observed view. Detailed dataset statistics and descriptions are provided in Appendix C.1 for reproducibility.

2) Compared Methods: We compare TIGER with eight IMVC baselines: **BRIDGE** (Jiang et al., 2025), **SCVT** (Dong et al., 2025), **FreeCSL** (Dai et al., 2025), **GAVIM** (Chen et al., 2026b), **IMVC-NCD** (Guo et al., 2026), **DGIMVCM** (Zhang et al., 2026b), **MVP** (Gao & Pu, 2025), and **CAMEL** (Gu et al., 2026). These methods represent advances in incomplete multi-view clustering from different perspectives. Detailed descriptions and implementation settings are provided in Appendix C.1.

3) Evaluation Metrics: Clustering quality is assessed using Clustering Accuracy (ACC), Normalized Mutual Information (NMI), and Adjusted Rand Index (ARI). ACC measures the agreement between predicted assignments and ground-truth labels after label matching, NMI quantifies the statistical dependence between the two partitions, and ARI evaluates pairwise partition consistency with chance agreement removed. Larger values of all three metrics indicate better clustering performance.

4.2 Experimental Results

1) Experimental Results Analysis: Table 1 compares TIGER with eight baselines on seven datasets under four missing rates. TIGER ranks first in all 84 evaluation settings, demonstrating consistent superiority across datasets and missing rates. On Synthetic3D at a missing rate of 0.1, TIGER achieves 96.50%, 85.58%, and 89.86% in ACC, NMI, and ARI, improving the strongest baseline by 0.36, 0.46, and 0.82 percentage points. Under severe missingness at 0.7, TIGER obtains 88.80%, 79.10%, and 79.25% on HW, exceeding the second best results by 1.35, 0.47, and 1.39 points. The advantage is most pronounced on ALOI-100, where TIGER surpasses the strongest competitors by 23.71, 13.84, and 6.27 points. These gains demonstrate the robustness of target guided retrieval, graph structure decoupling, and reliability guided intervention under severe incompleteness.

2) Ablation Study: We conduct ablation experiments on MNIST4, ALOI-100, and MNIST-Fea at a missing rate of 0.1 to assess each objective in TIGER. As shown in Table 2, the full objective achieves the best ACC, NMI, and ARI on all three datasets, indicating complementary effects among the components. Removing L_{str} substantially degrades performance on ALOI-100 and MNIST-Fea, confirming the importance of graph-structure decoupling for transferable cross-view relations. Removing L_{gii} causes a large drop on MNIST4, showing that intervention-guided target-view completion is critical under structural ambiguity. Although L_{rec} alone remains competitive on some datasets, jointly optimizing L_{rec}, L_{str}, and L_{gii} yields the most stable and accurate performance.

3) Parameter Analysis: Fig. 2 shows the sensitivity of TIGER to λ_{str} and λ_{gii} at a missing rate of 0.1. Broad high-ACC regions appear on all four datasets. HW and Synthetic3D remain nearly unchanged across the grid, while MNIST4 and MNIST-Fea show minor fluctuations. This indicates TIGER is robust to variations in the two loss weights and does not require finely tuned hyperparameters.

Table 1: Results (mean ± std, %) on seven datasets under four missing rates. Best and second-best results are highlighted in red and blue, respectively.

Dataset	Method	0.1			0.3			0.5			0.7		
		ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI
Synthetic3D	BRIDGE (ICCV'25)	95.50±0.83	82.32±2.46	87.04±2.29	90.90±1.23	72.51±2.02	75.76±2.21	86.96±0.42	62.93±2.06	66.65±1.97	83.58±0.92	55.64±1.27	57.80±1.64
	SCVT (TIP'25)	95.84±0.34	83.46±0.78	88.00±0.88	**92.16±0.29**	73.73±0.79	78.77±1.01	**87.81±0.95**	62.73±1.02	67.52±0.67	**83.59±0.41**	54.51±0.92	58.48±0.80
	FreeCSL (CVPR'25)	94.66±0.16	80.75±0.43	84.88±0.47	83.99±0.77	66.07±0.25	68.67±0.71	72.37±1.07	53.16±0.68	53.74±1.21	63.16±1.16	36.29±0.61	35.04±1.24
	GAVIM (TKDE'26)	80.25±2.25	54.26±2.25	54.82±3.47	77.79±3.90	50.08±3.29	51.23±4.97	74.17±5.84	45.72±4.60	46.68±6.00	71.92±6.58	41.52±5.06	42.21±7.19
	IMVC-NCD (ICML'26)	94.58±0.75	81.09±1.65	84.77±1.97	90.73±1.21	71.91±1.78	75.21±2.15	85.82±0.68	60.94±2.65	66.20±2.27	82.42±1.41	53.07±2.55	56.64±2.83
	DGIMVCM (AAAI'26)	**96.14±0.08**	**85.12±0.30**	**89.04±0.23**	92.03±0.51	**75.44±0.51**	**79.46±0.63**	86.62±0.02	62.93±0.69	**67.59±1.03**	83.00±0.33	**56.98±0.29**	**58.93±0.85**
	MVP (ICLR'25)	95.25±0.25	82.19±0.78	86.46±0.67	91.04±0.02	71.66±0.70	76.45±0.36	86.33±0.35	**62.94±0.68**	63.87±0.02	82.08±0.08	52.85±0.09	55.49±0.24
	CAMEL (ICLR'26)	87.24±0.75	77.79±0.37	76.39±0.82	84.94±0.41	70.80±0.44	68.79±0.60	82.63±0.53	62.31±0.60	60.21±0.34	80.71±0.31	54.72±0.73	52.63±0.54
	TIGER (Ours)	**96.50±0.17**	**85.58±0.65**	**89.86±0.48**	**93.17±0.00**	**75.61±0.10**	**80.80±0.05**	**88.00±1.00**	**63.27±2.62**	**67.81±2.12**	**84.25±0.25**	**57.41±0.17**	**59.64±0.93**
HW	BRIDGE (ICCV'25)	79.74±1.90	78.11±0.03	71.37±0.64	74.54±4.03	70.02±2.02	63.50±3.46	70.35±5.82	62.88±3.84	55.07±5.68	66.27±8.28	53.99±5.78	43.50±8.58
	SCVT (TIP'25)	81.12±1.52	70.99±0.85	64.01±2.53	73.12±1.02	63.36±0.83	55.59±1.57	65.72±0.84	56.75±0.34	45.52±1.24	57.40±0.31	48.08±0.37	35.96±0.28
	FreeCSL (CVPR'25)	78.33±1.12	80.61±1.38	70.62±0.63	76.12±1.10	77.37±1.41	66.31±1.17	73.67±0.81	70.87±1.15	63.35±1.47	71.72±0.63	67.32±0.68	59.16±1.79
	GAVIM (TKDE'26)	88.53±1.62	82.28±1.68	78.55±2.55	85.15±1.46	77.60±1.00	72.79±2.27	81.23±1.94	72.27±0.56	65.80±1.50	77.88±1.43	67.23±0.53	61.76±1.17
	IMVC-NCD (ICML'26)	95.47±1.09	92.34±0.75	**92.89±0.28**	92.13±1.02	**88.93±1.02**	**88.00±0.62**	86.36±1.72	83.84±1.59	**82.57±1.33**	80.34±1.33	**78.63±1.43**	**77.86±1.58**
	DGIMVCM (AAAI'26)	93.88±0.12	91.88±1.43	92.44±1.05	92.24±0.14	87.28±0.89	86.31±0.71	**90.18±0.47**	**84.15±1.06**	82.26±0.68	87.20±0.25	78.37±0.43	76.69±0.52
	MVP (ICLR'25)	95.08±0.38	89.78±0.60	89.42±0.78	81.92±1.59	81.04±1.52	69.37±2.69	68.00±2.69	67.75±2.31	53.09±5.29	56.08±4.18	54.60±3.35	35.18±6.96
	CAMEL (ICLR'26)	**95.53±0.07**	**93.76±0.20**	91.95±0.17	**93.06±0.02**	87.62±0.71	87.88±0.04	89.93±0.22	81.46±0.49	82.28±0.41	**87.45±0.25**	76.17±0.67	**77.59±0.55**
	TIGER (Ours)	**97.68±0.02**	**94.61±0.06**	**94.87±0.06**	**97.20±0.05**	**93.67±0.01**	**93.89±0.09**	**94.35±0.15**	**88.22±0.23**	**88.03±0.29**	**88.80±0.05**	**79.10±0.26**	**79.25±0.14**
MNIST4	BRIDGE (ICCV'25)	86.72±0.54	67.94±0.54	69.16±1.12	86.38±0.82	67.68±1.03	68.79±0.85	86.01±0.74	67.42±0.95	68.51±0.86	85.76±0.67	67.22±1.04	68.19±1.17
	SCVT (TIP'25)	88.36±0.79	70.25±0.65	70.44±0.74	87.62±0.94	69.71±0.89	69.95±0.80	86.75±0.55	69.16±1.21	69.28±0.94	86.13±0.22	68.37±0.58	68.47±0.52
	FreeCSL (CVPR'25)	87.73±1.21	73.17±0.68	74.39±0.72	86.98±1.12	72.13±0.30	73.78±0.52	86.02±0.43	70.95±0.92	72.80±1.10	85.26±0.14	70.14±0.48	71.88±0.31
	GAVIM (TKDE'26)	89.64±0.01	73.58±0.34	75.11±0.19	89.06±0.02	72.27±0.45	74.15±0.37	88.48±0.31	70.62±1.04	72.70±0.45	87.79±0.24	69.29±0.86	71.29±0.72
	IMVC-NCD (ICML'26)	89.24±0.27	**74.25±0.37**	**75.63±0.42**	88.91±0.16	**73.85±0.35**	**75.16±0.60**	**88.50±0.36**	**73.49±0.56**	**74.67±0.06**	**88.16±0.54**	**73.11±0.24**	**74.08±0.28**
	DGIMVCM (AAAI'26)	80.39±0.11	66.38±0.20	57.94±0.76	77.18±0.30	65.05±0.20	56.89±0.79	73.11±0.02	63.70±0.88	55.77±0.52	70.05±0.32	66.94±0.92	58.16±0.76
	MVP (ICLR'25)	88.80±0.08	71.44±0.11	73.42±0.11	87.58±0.04	69.68±0.17	70.86±0.09	86.52±0.16	68.05±0.53	68.88±0.63	85.76±0.16	66.64±0.88	67.28±0.30
	CAMEL (ICLR'26)	**90.16±0.24**	74.22±0.73	74.51±0.58	**89.28±0.35**	71.97±0.90	72.89±0.58	88.40±0.11	70.16±0.26	70.95±1.20	87.70±0.72	68.45±0.59	69.35±0.67
	TIGER (Ours)	**91.03±0.90**	**77.68±2.38**	**78.71±2.03**	**90.40±0.02**	**75.97±0.09**	**77.33±0.04**	**90.18±0.00**	**75.37±0.08**	**76.79±0.00**	**90.09±0.26**	**75.35±0.35**	**76.67±0.51**
CUB	BRIDGE (ICCV'25)	44.92±3.58	47.02±2.44	27.82±1.91	41.90±2.29	42.26±2.40	26.63±1.79	38.74±1.41	38.15±2.69	25.65±1.36	35.08±0.08	33.91±2.43	36.63±0.89
	SCVT (TIP'25)	60.00±1.33	62.88±1.99	44.84±1.66	54.94±1.81	56.90±2.41	44.08±1.31	50.46±2.35	51.16±2.16	43.23±1.16	45.09±2.41	44.55±1.99	46.45±1.12
	FreeCSL (CVPR'25)	63.09±1.91	59.05±0.46	45.59±0.13	53.40±1.70	50.91±1.21	40.91±0.31	44.21±1.17	45.57±2.01	37.27±0.65	37.00±0.83	37.36±2.38	33.46±1.28
	GAVIM (TKDE'26)	75.00±4.00	69.78±4.79	58.41±6.23	67.04±3.15	61.09±3.34	55.13±4.62	61.82±2.08	55.70±1.48	51.17±2.16	54.08±1.42	49.10±0.13	46.96±0.25
	IMVC-NCD (ICML'26)	60.17±1.50	57.50±0.47	45.67±1.71	57.03±1.83	53.04±0.79	44.46±1.00	52.98±1.81	49.32±0.68	43.16±0.51	50.09±2.42	45.12±1.11	49.88±0.25
	DGIMVCM (AAAI'26)	46.67±0.00	52.87±0.23	29.16±1.23	45.37±0.89	50.55±0.02	28.56±1.65	44.37±2.33	48.43±0.22	28.15±2.12	49.00±3.33	46.90±0.02	45.70±2.71
	MVP (ICLR'25)	**82.08±1.25**	77.08±2.70	**66.14±1.03**	**76.43±1.29**	69.48±2.39	**61.85±1.82**	67.73±1.56	60.04±1.93	**57.30±2.26**	60.33±1.33	52.72±1.74	**53.05±2.77**
	CAMEL (ICLR'26)	81.51±0.75	**79.58±0.36**	65.35±0.89	76.20±1.00	**76.90±0.42**	61.04±2.66	**69.94±1.01**	**75.36±0.84**	56.31±4.09	**65.12±0.73**	**74.08±0.43**	51.19±5.71
	TIGER (Ours)	**82.50±0.83**	**80.21±0.37**	**67.37±1.27**	**79.85±0.55**	**78.16±0.55**	**64.53±0.87**	**76.53±0.23**	**76.89±0.70**	**60.85±0.46**	**74.17±0.00**	**75.07±0.81**	**57.12±0.12**
WebKB	BRIDGE (ICCV'25)	79.94±0.38	36.15±0.50	35.07±0.86	76.75±1.63	31.83±1.34	32.50±1.00	72.22±2.21	25.88±2.07	28.79±0.65	68.63±3.28	20.10±2.82	24.68±0.98
	SCVT (TIP'25)	84.11±0.47	37.93±1.00	**41.41±0.96**	79.33±0.73	33.10±1.24	35.49±0.77	75.12±1.42	27.22±1.49	31.24±1.14	69.19±1.76	22.45±1.39	25.73±0.81
	FreeCSL (CVPR'25)	74.03±1.80	35.29±2.50	28.75±3.32	67.55±2.55	30.85±2.84	27.56±3.05	63.11±3.05	26.43±3.01	25.76±2.61	56.75±4.43	23.61±3.19	24.53±2.39
	GAVIM (TKDE'26)	62.71±0.34	32.18±1.33	22.74±1.05	58.25±0.71	29.32±1.27	21.77±1.27	53.04±0.02	26.63±1.63	21.06±1.55	47.17±0.41	24.01±1.74	28.38±1.93
	IMVC-NCD (ICML'26)	77.24±0.48	40.15±1.36	40.06±0.57	66.78±0.21	33.89±1.21	35.93±0.58	61.86±0.43	**27.78±0.53**	28.56±0.02	51.57±0.00	26.30±0.14	24.01±0.04
	DGIMVCM (AAAI'26)	**84.25±0.29**	**40.88±1.21**	40.16±1.32	**80.44±0.60**	**33.94±1.13**	**37.78±1.32**	**77.55±0.67**	27.46±0.89	**34.43±1.18**	**74.35±1.48**	26.78±1.30	**31.31±0.81**
	MVP (ICLR'25)	69.27±0.57	30.50±0.17	29.87±0.76	67.73±3.12	29.55±1.22	28.93±0.91	66.00±5.57	27.73±1.21	28.29±0.45	64.65±7.94	**28.17±2.17**	27.39±0.54
	CAMEL (ICLR'26)	80.54±0.28	40.77±0.36	40.72±0.35	77.21±0.65	33.11±0.02	34.19±0.61	72.73±0.77	27.75±0.26	28.72±0.64	69.31±1.33	21.59±0.27	20.08±0.61
	TIGER (Ours)	**85.30±0.05**	**42.12±0.14**	**48.34±0.11**	**81.92±1.14**	**34.52±1.67**	**39.54±2.75**	**80.16±1.09**	**27.94±1.46**	**34.92±2.46**	**79.69±1.47**	**28.35±2.35**	**34.04±3.33**
ALOI-100	BRIDGE (ICCV'25)	78.80±0.60	84.36±0.14	67.32±0.27	67.63±0.28	72.39±0.02	53.34±0.46	45.69±0.79	63.07±0.39	42.28±0.02	30.38±0.37	52.36±0.23	28.47±0.36
	SCVT (TIP'25)	77.14±0.24	83.11±0.77	66.73±0.62	63.93±0.62	75.91±0.66	**62.91±1.04**	55.72±0.86	66.06±0.93	**56.69±1.18**	**48.17±0.93**	60.55±0.47	49.87±0.81
	FreeCSL (CVPR'25)	77.85±0.61	84.19±0.07	66.63±0.54	60.41±1.32	75.89±0.45	55.15±0.87	45.90±1.47	66.93±1.37	41.84±1.30	31.09±2.22	57.88±1.44	29.93±1.66
	GAVIM (TKDE'26)	75.67±0.25	83.09±0.14	64.98±0.00	63.47±0.02	**76.41±0.02**	55.72±0.19	52.98±0.10	66.57±0.46	44.83±0.02	37.29±0.23	55.09±0.02	35.86±0.22
	IMVC-NCD (ICML'26)	74.96±0.48	82.14±0.35	67.17±0.37	66.67±0.05	75.74±0.24	61.10±0.28	**59.03±0.66**	**70.54±1.09**	54.81±0.41	47.72±0.24	**62.24±0.75**	**50.43±0.46**
	DGIMVCM (AAAI'26)	69.35±0.62	78.42±0.79	60.48±0.52	55.10±0.61	70.50±1.17	55.05±0.65	44.61±0.35	61.70±1.11	47.31±0.18	35.44±0.13	50.36±0.78	39.17±0.38
	MVP (ICLR'25)	**78.89±0.08**	**84.37±0.12**	**68.35±0.03**	64.64±0.22	74.98±0.26	59.11±0.18	50.26±0.73	64.39±0.24	42.91±0.39	39.46±0.71	53.68±0.15	29.71±0.28
	CAMEL (ICLR'26)	77.72±0.34	79.66±0.52	66.99±0.49	**69.08±0.41**	69.38±0.56	56.32±0.33	52.02±0.26	61.48±0.20	43.01±0.25	40.25±0.24	55.45±0.06	30.62±0.02
	TIGER (Ours)	**79.55±0.68**	**85.28±0.67**	**68.61±2.09**	**77.96±0.24**	**82.05±0.48**	**65.99±1.24**	**76.20±0.49**	**80.84±0.33**	**63.94±0.71**	**71.88±0.37**	**76.08±0.20**	**56.70±0.44**
MNIST-Fea	BRIDGE (ICCV'25)	96.42±0.14	94.31±0.26	94.31±0.37	94.08±0.38	91.75±0.46	91.05±0.38	90.67±1.43	88.84±0.33	89.32±0.98	87.12±1.41	86.71±0.28	86.18±0.63
	SCVT (TIP'25)	97.51±0.31	94.01±0.14	95.77±0.35	95.60±0.69	93.02±0.43	94.87±0.50	93.91±0.92	92.26±0.02	93.95±1.11	92.18±0.64	91.27±0.42	93.17±0.51
	FreeCSL (CVPR'25)	98.21±0.12	95.35±0.41	96.04±0.21	94.00±0.08	91.40±0.15	87.47±0.16	87.57±0.45	85.66±0.57	83.74±0.15	82.17±0.25	80.14±0.54	78.18±0.29
	GAVIM (TKDE'26)	97.87±0.26	94.64±0.81	95.38±0.74	96.10±0.68	93.34±0.97	94.95±0.88	93.80±0.38	92.20±0.37	94.54±0.92	91.28±0.68	90.85±0.81	94.07±0.73
	IMVC-NCD (ICML'26)	98.14±0.08	95.36±0.71	**96.42±0.34**	97.51±0.19	94.47±0.36	**95.85±0.05**	97.01±0.02	93.27±0.48	94.91±0.51	96.45±0.17	92.43±0.55	94.16±0.33
	DGIMVCM (AAAI'26)	98.11±0.25	95.07±0.30	96.08±0.04	96.90±0.15	94.13±0.27	95.63±0.21	95.55±0.69	93.49±0.18	95.17±0.70	94.52±0.26	92.43±0.42	94.75±0.18
	MVP (ICLR'25)	**98.39±0.00**	**95.45±0.04**	96.30±0.01	**97.62±0.02**	**94.94±0.07**	95.84±0.23	**97.32±0.37**	**94.26±0.02**	**95.43±0.02**	**97.16±0.04**	**93.70±0.09**	**94.97±0.08**
	CAMEL (ICLR'26)	96.75±0.14	86.98±0.65	88.52±0.48	93.44±0.21	86.27±0.43	87.73±0.56	89.53±0.02	85.28±0.19	86.66±0.78	85.81±0.34	84.24±0.32	85.40±0.32
	TIGER (Ours)	**98.48±0.00**	**95.68±0.02**	**96.68±0.01**	**98.37±0.03**	**95.37±0.06**	**96.43±0.06**	**98.21±0.01**	**94.95±0.02**	**96.09±0.02**	**97.99±0.01**	**94.42±0.03**	**95.61±0.03**

Table 2: Ablation results (mean ± std, %) on MNIST4, ALOI-100, and MNIST-Fea at a missing rate of 0.1 (10%). Best and second-best results are highlighted in red and blue, respectively.

Losses			MNIST4			ALOI-100			MNIST-Fea		
L_{rec}	L_{str}	L_{gii}	ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI
✓	×	×	78.51 ± 5.44	55.45 ± 5.48	55.04 ± 8.15	**75.50 ± 0.40**	**82.82 ± 0.47**	**64.46 ± 0.62**	97.66 ± 0.08	93.53 ± 0.24	94.89 ± 0.18
✓	×	✓	**87.19 ± 0.51**	**68.00 ± 0.38**	**70.25 ± 1.11**	38.72 ± 1.87	60.48 ± 0.87	28.08 ± 1.67	86.73 ± 0.03	88.87 ± 0.09	83.44 ± 0.12
✓	✓	×	52.58 ± 7.52	31.28 ± 5.07	29.49 ± 6.66	43.22 ± 0.44	62.87 ± 0.32	31.06 ± 0.20	**97.86 ± 0.07**	**94.25 ± 0.19**	**95.32 ± 0.18**
✓	✓	✓	**91.03 ± 0.90**	**77.68 ± 2.38**	**78.71 ± 2.03**	**79.55 ± 0.68**	**85.28 ± 0.67**	**68.61 ± 2.09**	**98.48 ± 0.00**	**95.68 ± 0.02**	**96.68 ± 0.01**

4) Convergence Analysis: Fig. 3 compares the optimization trajectories of Ours and the competing methods on four datasets at a missing rate of 0.1. In these curves, Ours reaches a stable plateau earlier, typically around 70–80 epochs, whereas most methods stabilize later, roughly between 90 and 110 epochs. Notably, several baselines start from comparable or even higher ACC values on some datasets, but Ours catches up rapidly and attains the strongest final performance. This behavior reflects faster optimization and stable late-stage convergence under incomplete-view settings.

5) Representation Visualization: Fig. 4 shows the evolution of the learned representation on ALOI-100. As training proceeds from 25 to 100 epochs, the initially mixed and scattered samples gradually form more compact and clearly separated clusters, while ACC increases from 45.18% to 80.23%. This progressive separation is consistent with the convergence trend in Fig. 3, indicating that TIGER steadily learns increasingly discriminative representations during optimization.

6) t-SNE Comparison: Fig. 5 compares the learned representations of CAMEL, DGIMVCM, IMVC-NCD, and TIGER on MNIST-Fea at a missing rate of 0.3. TIGER forms more compact and clearly separated clusters with less inter-cluster mixing across different semantic groups. This is consistent

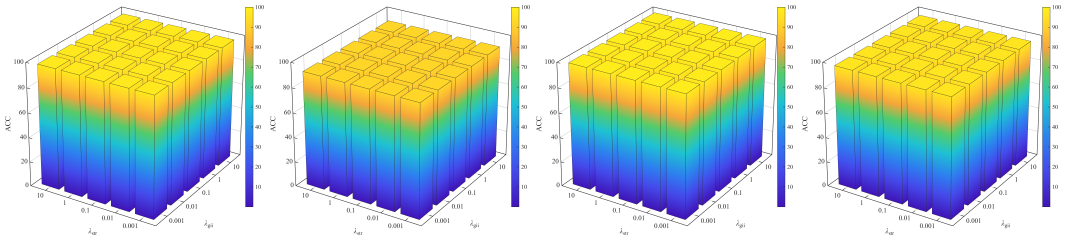

(a) HW (b) MNIST4 (c) MNIST-Fea (d) Synthetic3D

Figure 2: Parameter sensitivity with respect to λ_{str} and λ_{gii} on four datasets at a missing rate of 0.1.

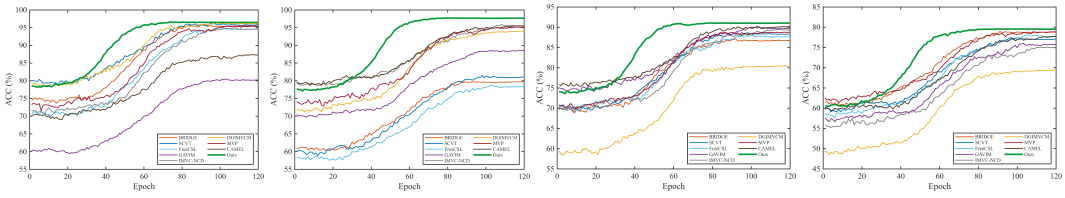

(a) Synthetic3D (b) HW (c) MNIST4 (d) ALOI-100

Figure 3: Illustrative convergence curves on four datasets at a missing rate of 0.1.

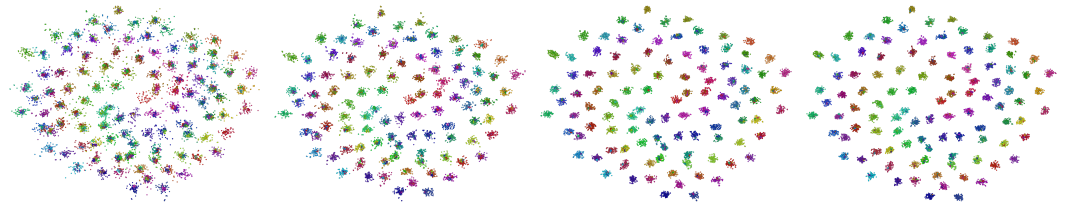

(a) 25 epoch (45.18%) (b) 50 epoch (63.75%) (c) 75 epoch (78.91%) (d) 100 epoch (80.23%)

Figure 4: t-SNE visualization of representations on ALOI-100 at different training epochs.



(a) CAMEL (b) DGIMVCM (c) IMVC-NCD (d) TIGER (Ours)

Figure 5: t-SNE visualization on MNIST-Fea at a 30% missing rate.

with Table 1, where TIGER achieves the best ACC under 30% missingness, further supporting the effectiveness of the proposed target-aware retrieval and intervention mechanisms.

5 Conclusion

In this paper, we proposed TIGER, a target aware interventional graph structure decoupling and retrieval framework for incomplete multi view clustering. TIGER separates stable structure from view dependent variations, enabling cross view relations to guide missing view recovery. Instead of reconstructing missing features, TIGER retrieves anchors from observed target samples using structural evidence. We further introduced confidence guided graph intervention and residual completion to suppress unreliable local dependencies, while reliability aware fusion limits the influence of uncertain recovered representations. Extensive experiments on seven benchmark datasets demonstrate that TIGER consistently outperforms strong recent baselines. These results show that combining target view geometry with reliable structural transfer provides an effective solution for IMVC.

REFERENCES

Guoqing Chao, Yi Jiang, and Dianhui Chu. Incomplete contrastive multi-view clustering with high-confidence guiding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 11221–11229, 2024.

Guoqing Chao, Zhenghao Zhang, Lei Meng, Jie Wen, and Dianhui Chu. Federated incomplete multi-view clustering with globally fused graph guidance. In *Proceedings of the 42nd International Conference on Machine Learning*, 2025.

Jingyi Chen, Zhibin Dong, Tiejun Li, and Yibo Han. DC-SPAN: A dual contrastive attention network for multi-view clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 20127–20135, 2026a.

Wenlan Chen, Lu Gao, Cheng Liang, and Fei Guo. Deep variational incomplete multi-view clustering with information-theoretic guidance. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pp. 2457–2466, 2025.

Wenlan Chen, Lu Gao, Daoyuan Wang, Fei Guo, and Cheng Liang. Geometry-aware variational information maximization for deep incomplete multi-view clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 20289–20297, 2026b.

Wenlan Chen, Lu Gao, Daoyuan Wang, Cheng Liang, and Fei Guo. Information-theoretic disentangled latent modeling with conditional diffusion for incomplete multi-view clustering. In *Proceedings of the 43rd International Conference on Machine Learning*, 2026c.

Yuzhuo Dai, Jiaqi Jin, Zhibin Dong, Siwei Wang, Xinwang Liu, En Zhu, Xihong Yang, Xinbiao Gan, and Yu Feng. Imputation-free and alignment-free: Incomplete multi-view clustering driven by consensus semantic learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5071–5081, 2025.

Xiaojian Ding, Lin Zhao, Xian Li, and Xiaoying Zhu. Incomplete multi-view clustering via hierarchical semantic alignment and cooperative completion. In *Advances in Neural Information Processing Systems*, pp. 52175–52195, 2025.

Shihao Dong, Yue Liu, Xiaotong Zhou, Yuhui Zheng, Huiying Xu, and Xinzhong Zhu. RAC-DMVC: Reliability-aware contrastive deep multi-view clustering under multi-source noise. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 20835–20843, 2026.

Zhibin Dong, Dayu Hu, Jiaqi Jin, Siwei Wang, Xinwang Liu, and En Zhu. Selective cross-view topology for deep incomplete multi-view clustering. *IEEE Transactions on Image Processing*, 34:4792–4805, 2025.

Yiming Du, Ziyu Wang, Jian Li, Rui Ning, and Lusi Li. Deep incomplete multi-view clustering via hierarchical imputation and alignment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 20941–20949, 2026.

Ran Eisenberg, Jonathan Svirsky, and Ofir Lindenbaum. COPER: Correlation-based permutations for multi-view clustering. In *International Conference on Learning Representations*, pp. 46154–46179, 2025.

Wei Feng, Danting Liu, Qianqian Wang, Mengping Jiang, and Bin Liu. Federated incomplete multi-view clustering with tensorized low-rank constraint. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 21083–21091, 2026.

Xin Gao and Jian Pu. Deep incomplete multi-view learning via cyclic permutation of vaes. In *International Conference on Learning Representations*, pp. 99609–99643, 2025.

Zhibin Gu, Zhenhao Zhong, Xi Zhang, and Bing Li. Constant degree matrix-driven incomplete multi-view clustering via connectivity-structure and embedding tensor learning. In *International Conference on Learning Representations*, pp. 84649–84671, 2026.

Qian Guo, Gaohui Zuo, Bingbing Jiang, Guangrui Fan, Zhihua Cui, Xinyan Liang, and Jianjian Ding. Incomplete multi-view clustering via neighborhood-conditioned diffusion. In *Proceedings of the 43rd International Conference on Machine Learning*, 2026.

Xiaorui Jiang, Buyun He, Peng Yuan Zhou, Xinyue Chen, Jingcai Guo, Jie Xu, and Yong Liao. A unified framework to bridge complete and incomplete deep multi-view clustering under non-iid missing patterns. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 594–603, 2025.

Jiaqi Jin, Siwei Wang, Zhibin Dong, Xinwang Liu, and En Zhu. 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*, pp. 11600–11609, 2023.

Wenyuan Kong, Zhibin Gu, and Bing Li. Dual-branch representations with dynamic gated fusion and triple-granularity alignment for deep multi-view clustering. In *International Conference on Learning Representations*, 2026.

Jiyuan Liu, Xinwang Liu, Xinhang Wan, Ke Liang, Weixuan Liang, Sihang Zhou, Huijun Wu, and Kehua Guo. Incomplete multi-view deep clustering with data imputation and alignment. In *Advances in Neural Information Processing Systems*, pp. 89119–89142, 2025.

Yanxi Liu, Fangxi Liu, Yipin Hu, Yanwei Yu, Lei Meng, Yongyong Chen, and Guoqing Chao. Dual-stage contrastive learning-enhanced multi-view variational clustering. In *Proceedings of the 43rd International Conference on Machine Learning*, 2026.

Shuangmei Peng, Zhe Chen, Tianyang Xu, and Xiaojun Wu. Relationship alignment for view-aware multi-view clustering. In *International Conference on Learning Representations*, pp. 147335–147347, 2026.

Yalan Qin, Guorui Feng, and Xinpeng Zhang. Fast incomplete multi-view clustering by flexible anchor learning. In *Proceedings of the 42nd International Conference on Machine Learning*, 2025.

Fujian Ren, Wenlan Chen, Lu Gao, Fei Guo, and Cheng Liang. Dual-level distribution alignment for deep incomplete multi-view clustering. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pp. 2476–2485, 2025.

Bocheng Wang, Chusheng Zeng, Mulin Chen, and Xuelong Li. Towards learnable anchor for deep multi-view clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 21044–21052, 2025.

Jing Wang, Songhe Feng, Gengyu Lyu, and Jiazheng Yuan. SURER: Structure-adaptive unified graph neural network for multi-view clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 15520–15527, 2024.

Qianqian Wang, Xurui Liao, Wei Feng, and Quanxue Gao. Debiased and denoised representation learning for incomplete multi-view clustering. In *International Conference on Learning Representations*, pp. 73338–73350, 2026a.

Qianqian Wang, Haiming Xu, Wei Feng, and Quanxue Gao. Adversarial fair incomplete multi-view clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 10011–10019, 2026b.

Xinxin Wang, Yongshan Zhang, Xiaochen Yuan, and Yicong Zhou. Cross-view anchor graph learning and factorization for incomplete multi-view clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 26570–26578, 2026c.

Bingjun Wei, Xuemei Cao, Jiafen Liu, Haoyang Liang, and Xin Yang. Enhanced federated deep multi-view clustering under uncertainty scenario. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 26806–26813, 2026.

Jie Wen, Chengliang Liu, Gehui Xu, Zhihao Wu, Chao Huang, Lunke Fei, and Yong Xu. Highly confident local structure based consensus graph learning for incomplete multi-view clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15712–15721, 2023.

Jie Wen, Shijie Deng, Waikeung Wong, Guoqing Chao, Chao Huang, Lunke Fei, and Yong Xu. Diffusion-based missing-view generation with the application on incomplete multi-view clustering. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.

Baili Xiao, Ke Liang, Jiaqi Jin, Jun Wang, Yinbo Xu, Siwei Wang, and En Zhu. Dmcar: Disentangled mixture-of-experts with context-aware routing for multi-view clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 27055–27063, 2026.

Deng Xu, Chao Zhang, Cong Guo, Chunlin Chen, and Huaxiong Li. Fast incomplete multi-view clustering with adaptive similarity completion and reconstruction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 21734–21742, 2025.

Xiaoqiang Yan, Zhixiang Jin, Fengshou Han, and Yangdong Ye. Differentiable information bottleneck for deterministic multi-view clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 27435–27434, 2024.

Xihong Yang, Siwei Wang, Fangdi Wang, Jiaqi Jin, Suyuan Liu, Yue Liu, En Zhu, Xinwang Liu, and Yueming Jin. Automatically identify and rectify: Robust deep contrastive multi-view clustering in noisy scenarios. In *Proceedings of the 42nd International Conference on Machine Learning*, 2025.

Shengju Yu, Zhibin Dong, Siwei Wang, Xinhang Wan, Yue Liu, Weixuan Liang, Pei Zhang, Wenxuan Tu, and Xinwang Liu. Towards resource-friendly, extensible and stable incomplete multi-view clustering. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.

Teng Zhang, Pengyuan Li, Zisen Kong, Dongxia Chang, and Yao Zhao. Bipartite contrastive multi-view clustering with singular value modulation. *Neurocomputing*, 656:131528–13139, 2025a.

Yuanyang Zhang, Yijie Lin, Weiqing Yan, Li Yao, Xinhang Wan, Guangyuan Li, Chao Zhang, Guanzhou Ke, and Jie Xu. Incomplete multi-view clustering via diffusion contrastive generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 22650–22658, 2025b.

Yuanyang Zhang, Xinhang Wan, Chao Zhang, Jie Xu, Cunjian Chen, Tien-Tsin Wong, Li Yao, and Yijie Lin. Mamba-driven multi-view discriminative clustering via global-local cross-view sequence modeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 28519–28527, 2026a.

Yunhe Zhang, Jinyu Cai, Zhihao Wu, Pengyang Wang, and See-Kiong Ng. Mixture of experts as representation learner for deep multi-view clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 22704–22713, 2025c.

Zhenghao Zhang, Jun Xie, Xingchen Chen, Tao Yu, Hongzhu Yi, Kaixin Xu, Yuanxiang Wang, Tianyu Zong, Xinming Wang, Jiahuan Chen, Guoqing Chao, Feng Chen, Zhepeng Wang, and Jungang Xu. Dynamic deep graph learning for incomplete multi-view clustering with masked graph reconstruction loss. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 28600–28608, 2026b.

Bowen Zhao, Qianqian Wang, Zhengming Ding, and Quanxue Gao. Attribute-missing multi-view graph clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 25832–25841, 2025.

Yunxiao Zhao, Liang Bai, and Xian Yang. Adaptive evolutionary fusion for multi-view clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 28769–28777, 2026.

Haochen Zhou, Guofeng Ding, Mouxing Yang, Peng Hu, Yijie Lin, and Xi Peng. Uncover underlying correspondence for robust multi-view clustering. In *International Conference on Learning Representations*, 2026.

Zhengzhong Zhu, Pei Zhou, Lanxi Bai, Jia Nie, Li Cheng, Shiquan Min, and Jiangping Zhu. Prototype and sample level semantic alignment for incomplete multi-view clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5818–5827, 2026.