MULTI-GRANULARITY RELATION TRANSFER WITH SHARED ANCHORS FOR INCOMPLETE MULTI-VIEW CLUSTERING

Anonymous authors
Paper under double-blind review

ABSTRACT

Multi-granularity modeling aims to capture complementary cluster structures from fine to coarse resolutions to enhance incomplete multi-view clustering. Unfortunately, view incompleteness disrupts such modeling by inducing two major problems: cross-view reference inconsistency and cross-granularity structural instability. The first is **reference inconsistency**, where missing views make relational structures depend on different observed sample subsets, making structures from different views and granularities difficult to integrate. Considering this, we introduce shared anchors as a unified relational reference and construct masked sample-anchor assignments, which are further aggregated into a sample-complete global reference without imputing missing views. The second is **multi-granularity instability**: fine-grained structures are sensitive to missingness and noise, while coarse-grained structures may suffer from semantic drift. Therefore, we develop an anchor-induced multi-granularity relation transfer mechanism to extract fine-to-coarse structures on compact anchor graphs and transfer them to samples. By integrating sparse local geometry, our model preserves supported neighborhood evidence and promotes local consistency. Additionally, adaptive low-rank consensus learning adjusts the relative view-granularity contributions according to their consistency with the evolving consensus. Extensive experiments conducted on multiple benchmark incomplete multi-view datasets demonstrate the superiority of the proposed method over existing state-of-the-art approaches.

1 INTRODUCTION

With the rapid development of multimedia and sensor technologies, real-world objects are increasingly represented by multiple heterogeneous views, such as the textual and visual content of a web page or different feature descriptors of an image (Fu et al. (2020); Qin et al. (2025b)). Since each view usually captures only partial characteristics of an object, integrating complementary information across views can provide a more comprehensive description of the underlying data structure. Multi-view clustering (MVC) therefore aims to exploit multiple views jointly to uncover their intrinsic cluster structure (Liu et al. (2026); Yang et al. (2026); Ji & Feng (2025); Long et al. (2025)). In practical applications, however, complete observations across all views are often unavailable because of acquisition failures, storage loss, or collection constraints. For instance, in medical analysis, a patient may be characterized by medical images, laboratory tests, and clinical records, while some examinations may be unavailable due to acquisition cost or practical constraints. This results in incomplete multi-view data, where some samples are missing one or more views. Incomplete multi-view clustering (IMVC) aims to partition such incomplete samples into meaningful clusters by jointly utilizing their available views (Wang et al. (2025a); Qin et al. (2025a); Zhao et al. (2025b)).

To handle incomplete multi-view data, existing IMVC methods have explored various strategies to exploit the information retained in the observed views. A straightforward line of research focuses on completion, where missing views or their latent representations are estimated from the available data before performing clustering (Lin et al. (2021); Wen et al. (2024); Wang et al. (2026b)). Another widely studied direction learns a shared latent representation by projecting different views into a common space, so that cross-view information can be integrated despite partial observations (Zeng et al. (2024); Wang et al. (2026a); Dai et al. (2025)). Graph-based methods characterize the

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.

relationships among samples by constructing view-specific or consensus graphs, enabling structural information to be propagated across incomplete views (Chao et al. (2025b); Zhao et al. (2026); Ding et al. (2026)). To improve scalability, anchor-based methods replace dense sample-wise relationships with compact sample-anchor associations while retaining the main structural information of the data (Chen et al. (2024); He et al. (2025); Ji et al. (2026)). Tensor-based approaches, in contrast, organize view-specific representations or graphs into high-order tensors to capture correlations across multiple views (Li et al. (2022); Zhang et al. (2024); Gu et al. (2026)). Together, these methods provide mechanisms for recovering missing information, integrating cross-view representations, and preserving the structural relationships underlying incomplete multi-view data during clustering.

Existing IMVC methods mainly recover missing information or integrate the available views to obtain a unified representation or structure for clustering. However, describing the data with a single structural resolution may overlook clustering patterns that appear at different levels (Yu et al. (2025); Wang et al. (2025b)). Fine-grained structures capture detailed neighborhood relations and local variations, whereas coarse-grained structures reveal broader grouping patterns and overall cluster organization. Modeling these structures across multiple granularities can therefore provide complementary information for characterizing the underlying cluster structure (Yang et al. (2025); Wei et al. (2026)). Nevertheless, view incompleteness introduces difficulties when such multi-granularity structures are learned from different views. First, different views are observed on different subsets of samples, causing their relational representations to be constructed from incomplete observations. As a result, structures extracted from different views and granularities lack a unified reference for effective integration and comparison (Zhao et al. (2026)). Second, structural reliability varies across granularities under incomplete observations. Fine-grained structures are more sensitive to missing samples, noise, and neighborhood perturbations, while coarse-grained structures may lose discriminative details or distort the underlying cluster organization through excessive aggregation. These difficulties hinder the construction of comparable and reliable multi-granularity relations for IMVC.

To address these challenges, we propose an *Anchor-induced Multi-Granularity Consensus framework* (AMGC) for incomplete multi-view clustering. We first introduce shared anchors to establish a common relational reference across views and construct masked sample-anchor assignments using only the available observations. These assignments are further aggregated into a global reference covering all samples, providing aligned relational coordinates for subsequent structure modeling without imputing missing views. Based on the shared anchor space, AMGC constructs compact anchor graphs at multiple resolutions and transfers the resulting fine-to-coarse structures back to the sample level, thereby preserving complementary clustering information across granularities. Sparse local geometry helps retain neighborhood-supported relations and preserve local structural details during consensus learning. Finally, an adaptive low-rank consensus model jointly integrates the resulting view-granularity relations and adjusts their relative contributions according to their consistency with the evolving consensus. The low-rank formulation also enables efficient matrix-free optimization of the global relations. The main contributions of this work are summarized as follows:

- We develop a shared-anchor relational reference mechanism for incomplete multi-view clustering, which establishes an anchor-based reference space across partially observed views and constructs a sample-complete global reference without imputing missing data.
- We develop an anchor-induced multi-granularity relation transfer mechanism to capture fine-to-coarse cluster structures and integrate them through adaptive low-rank consensus with sparse local geometry.
- We analyze the conditional stability of global relation transfer and its weighted fitting target, and derive matrix-free gradients for consensus learning. Extensive experiments on multiple benchmark datasets demonstrate the effectiveness and competitiveness of AMGC.

2 Related Work

2.1 Incomplete Multi-View Clustering

Recent IMVC studies have developed diverse techniques to recover cross-view information and learn clustering structures from partial observations. Early methods mainly exploit matrix factorization, shared embeddings, and graph learning. Hu & Chen (2019) jointly learn latent factors from partially observed views, while Wen et al. (2019) infer missing views through unified embedding alignment.

Graph-oriented approaches preserve sample relationships during clustering; adaptive graph learning (Wen et al. (2020)) constructs neighborhood structures from available samples, and subsequent graph refinement and completion methods further recover structural information across incomplete views (Li et al. (2024); Zhao et al. (2023)). With the development of deep learning, Completer (Lin et al. (2021)) predicts cross-view representations through contrastive learning, while dual contrastive prediction (Lin et al. (2023)) further exploits cross-view consistency in the latent space. Prototype-based methods establish sample–prototype correspondences to handle missing observations (Jin et al. (2023)), and graph-propagation methods transfer available information through latent neighborhood structures (Pu et al. (2024); Chao et al. (2024; 2025a)). More recent studies explore consensus-semantic learning without explicit data completion (Dai et al. (2025)) and refine distribution bias introduced by missing views (Wang et al. (2026a)). These studies have further advanced IMVC through improved information recovery, representation, and structural modeling techniques.

2.2 Multi-Granularity and Anchor-Based Structure Learning

Multi-granularity structure learning has recently attracted increasing attention in multi-view clustering because cluster organization can be characterized at different structural resolutions. MASTER (Wang et al. (2025b)) models invariant structures at sample, neighborhood, and category granularities, while MGE (Yang et al. (2025)) generates hierarchical partitions and integrates cross-view and cross-granularity information through co-association learning. Multi-scale graph methods capture structural information at different resolutions through graph diffusion and adaptive fusion (Chen et al. (2026)). Anchor learning provides another efficient way to represent such structures. DVSAI (Yu et al. (2024)) constructs view-shared anchors with different dimensions and sizes and integrates the multi-scale anchor graphs for incomplete multi-view clustering. Recent view-adaptive multi-granularity anchor learning (Wei et al. (2026)) learns anchor structures at different granularities according to individual views, while structure-aware anchor graph methods jointly optimize anchors and graph relations to preserve clustering structure (Guo et al. (2026)). However, varying anchor dimensions or counts, as in DVSAI, does not explicitly model fine-to-coarse grouping relations, while MGE's sample-partition ensemble does not explicitly handle missing views. AMGC addresses this gap through an explicit multi-granularity relation transfer mechanism, enabling grouping evidence at different resolutions to jointly better guide the clustering process under incomplete observations.

3 Method

The overall framework of AMGC is illustrated in Figure 1, including shared-anchor relational reference construction, multi-granularity relation transfer, and adaptive low-rank consensus learning. To describe the proposed framework clearly, we first introduce the notations and problem setting.

Notations. Let $\mathcal{X} = \{X^{(v)}\}_{v=1}^{V}$ denote an incomplete multi-view dataset with N samples and V views, where $X^{(v)} \in \mathbb{R}^{N \times d_v}$ represents the data matrix of the v-th view and d_v denotes its feature dimension. We use $M \in \{0,1\}^{N \times V}$ to indicate view availability, where $M_{iv} = 1$ if sample i is observed in view v, and $M_{iv} = 0$ otherwise. Each sample is assumed to contain at least one observed view, i.e., $\sum_{v=1}^{V} M_{iv} \geq 1$. Let $\Omega_v = \{i \mid M_{iv} = 1\}$ and $n_v = |\Omega_v|$ denote the observed sample set and its size in the v-th view, respectively. We consider $2 \leq K \leq N$ and $n_v \geq 2$ for $v = 1, \ldots, V$, where K is the number of clusters. The task is to partition these N samples with potentially missing views into K clusters under the given view-availability pattern for clustering.

3.1 Shared-Anchor Reference Construction

The observed sample sets $\{\Omega_v\}$ generally differ across incomplete views, resulting in heterogeneous sample-anchor relations that require view-specific observation patterns for effective aggregation. To establish a consistent relational reference, we select a set of data anchors whose sample identities are shared across all views. Each anchor corresponds to an actual sample index, so its identity remains fully fixed across views even when the corresponding observation is unavailable in a particular view.

Shared Anchor Selection. We first measure the observation completeness of sample i by $c_i = \sum_{v=1}^{V} M_{iv}$. Given an initial anchor number $m_0 \in \{2, \ldots, N\}$ and an expansion factor $\kappa > 1$, the $\min(\lceil \kappa m_0 \rceil, N)$ samples with the largest completeness scores form the initial candidate pool $\mathcal{C}$.

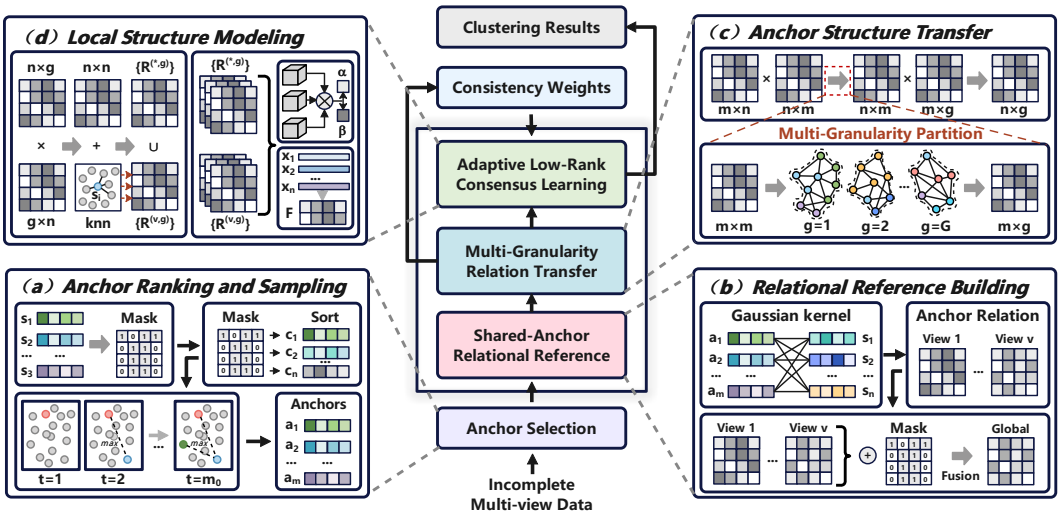

Figure 1: Overview of AMGC. Multi-granularity partitions are learned in a shared anchor space and transferred to sample relations through sample-anchor assignments. Local geometry refines view-specific relations, and adaptive low-rank consensus learning integrates the local and global relations.

This screening favors candidates observed in more views. Each view is subsequently standardized using its respective observed samples. Let $\mathcal{V}_{ij} = \{v : M_{iv}M_{jv} = 1\}$ denote the views jointly observed for samples i and j. For $\mathcal{V}_{ij} \neq \varnothing$, their dimension-normalized dissimilarity is defined as

$$\delta_{ij} = \frac{1}{|\mathcal{V}_{ij}|} \sum_{v \in \mathcal{V}_{ij}} \frac{\|\widetilde{x}_i^{(v)} - \widetilde{x}_j^{(v)}\|_2^2}{d_v}. \tag{1}$$

Starting from the candidate with the highest completeness score, we apply farthest-point sampling (FPS Gonzalez (1985)) to iteratively select the remaining candidate farthest from its nearest anchor:

$$a_t = \arg \max_{i \in \mathcal{C} \setminus \mathcal{A}_{t-1}} \min_{\substack{a \in \mathcal{A}_{t-1} \\ \mathcal{V}_{ia} \neq \varnothing}} \delta_{ia}, \quad t = 2, \ldots, m_0, \tag{2}$$

where $\mathcal{A}_{t-1} = \{a_1, \ldots, a_{t-1}\}$. We explicitly set $\min \varnothing = +\infty$ as a selection priority for candidates sharing no observed view with the selected anchors, rather than treating their unknown dissimilarities as measured distances. Selection ties are resolved by completeness, followed by random choice with a fixed seed. After m_0 selections, we further check the available views sequentially. If a view has no observed anchor, we add an unselected sample with the highest completeness score among its remaining observed samples. The final shared anchor set $\mathcal{A} = \{a_1, \ldots, a_m\}$, where $m = |\mathcal{A}| \geq m_0$, provides a consistent common reference through identical shared sample indices across views.

Sample-Anchor Assignment Construction. For each view v, we compute the relations between samples observed in that view and all anchors also observed in the same view. Specifically, let $\mathcal{D}_v^+$ denote the set of positive values of $q_{ij}^{(v)}$ for $i \in \Omega_v$ and $j \in \mathcal{I}_v$. They are defined as follows:

$$q_{ij}^{(v)} = \|\widetilde{x}_i^{(v)} - \widetilde{x}_{a_j}^{(v)}\|_2^2, \quad \sigma_v^2 = \begin{cases} \max\{\mathrm{med}_{\mathrm{low}}(\mathcal{D}_v^+), \epsilon\}, & \mathcal{D}_v^+ \neq \emptyset, \\ 1, & \text{otherwise.} \end{cases} \tag{3}$$

Here, $\mathcal{D}_v^+$ contains all positive $q_{ij}^{(v)}$ with $i \in \Omega_v$ and $j \in \mathcal{I}_v$. The operator $\mathrm{med}_{\mathrm{low}}$ returns the lower median, and $\sigma_v^2 = 1$ when $\mathcal{D}_v^+$ is empty. The Gaussian affinity for $i \in \Omega_v$ and $j \in \mathcal{I}_v$ is defined as

$$W_{ij}^{(v)} = \exp\left(-\frac{\|\widetilde{x}_i^{(v)} - \widetilde{x}_{a_j}^{(v)}\|_2^2}{2\sigma_v^2}\right). \tag{4}$$

The Gaussian affinities are normalized over the available anchors to obtain the sample-anchor assignment matrix $Z^{(v)} \in \mathbb{R}^{N \times m}$. Each observed sample is assigned to the anchors available in its view based on the normalized affinities, while unavailable sample-anchor pairs are masked out.

Global Relational Reference Construction. Since different views contain different observed sample-anchor pairs, we aggregate the view-specific assignments $\{Z^{(v)}\}_{v=1}^{V}$ under the shared anchor space. Specifically, we average over views in which both sample i and anchor a_j are observed as $\bar{Z}_{ij} = \frac{\sum_{v=1}^{V} M_{iv} M_{a_j v} Z_{ij}^{(v)}}{\sum_{v=1}^{V} M_{iv} M_{a_j v} + \epsilon}$, where $\epsilon > 0$ prevents division by zero and unsupported pairs receive zero values. The aggregated matrix $\bar{Z}$ is further row-normalized to obtain the global relational reference $Z^{(*)} \in \mathbb{R}^{N \times m}$ for subsequent multi-granularity structure learning.

3.2 Anchor-Based Multi-Granularity Structure Learning

Based on the global relational reference $Z^{(*)}$ and view-specific assignments ${Z^{(v)}}_{v=1}^{V}$, we further construct multi-granularity structures in the anchor space and transfer them to the sample level.

Anchor Graph Construction. The sample-anchor assignments provide a compact representation of sample relations in the anchor space. Let $\mathcal{S} = \{1, \ldots, V, *\}$ denote the set of relational sources, including the V view-specific assignments and the global reference $*$. For each relational source $s \in \mathcal{S}$, we construct an anchor graph by measuring the shared assignment patterns among anchors:

$$A_{\text{anc}}^{(s)} = (Z^{(s)})^{\top} Z^{(s)} \in \mathbb{R}^{m \times m}, \tag{5}$$

where larger entries indicate stronger correlations between anchors through shared sample assignments. Since local sources may contain unobserved anchors, we retain only observed anchors in each source. Specifically, let $\mathcal{I}_* = \{1, \ldots, m\}$ for the global source and $\mathcal{I}_v = \{j \mid M_{a_j v} = 1\}$ for a local view $s = v$. For each source $s \in \mathcal{S}$, let $m_s = |\mathcal{I}_s|$ denote the number of available anchors. The corresponding anchor graph and normalized form are defined as

$$A_{\text{act}}^{(s)} = A_{\text{anc}}^{(s)}[\mathcal{I}_s, \mathcal{I}_s], \quad \widehat{A}_{\text{act}}^{(s)} = (\Delta_{\text{act}}^{(s)})^{-1/2} A_{\text{act}}^{(s)} (\Delta_{\text{act}}^{(s)})^{-1/2}, \quad \Delta_{\text{act}}^{(s)} = \text{diag}(A_{\text{act}}^{(s)} \mathbf{1}). \tag{6}$$

The normalized graphs provide the structural basis for partitioning anchors at different granularities.

Multi-Granularity Anchor Partitioning. To capture structures at different resolutions, we perform spectral partitioning on normalized anchor graphs. Let $1 \leq c_1 < c_2 < \cdots < c_G \leq \min_{s \in \mathcal{S}} m_s$ denote the predefined anchor granularities, where larger c_g corresponds to finer partitions. For each relational source $s \in \mathcal{S}$ and granularity g, the observed anchor graph is partitioned as

$$Y_{\text{act}}^{(s,g)} = \text{SC}\Big(\widehat{A}_{\text{act}}^{(s)}, c_g\Big) \in \{0,1\}^{m_s \times c_g}, \quad Y_{\text{act}}^{(s,g)} \mathbf{1}_{c_g} = \mathbf{1}_{m_s}. \tag{7}$$

The anchor partition is extended to all anchors, with unavailable anchors assigned zero vectors:

$$Y_{\mathcal{I}_s,:}^{(s,g)} = Y_{\text{act}}^{(s,g)}, \quad Y_{\mathcal{I}_s^c,:}^{(s,g)} = 0, \quad Y^{(s,g)} \in \{0,1\}^{m \times c_g}. \tag{8}$$

The resulting $Y^{(s,g)}$ provides the multi-granularity anchor partition for each relational source.

Anchor-to-Sample Relation Transfer. The obtained anchor partitions are transferred to the sample space through the sample-anchor assignments. For each relational source $s \in \mathcal{S}$ and granularity g, the sample-level membership matrix according to the learned anchor groups is computed as

$$H^{(s,g)} = \text{RowNorm}\Big(Z^{(s)} Y^{(s,g)}\Big) \in \mathbb{R}^{N \times c_g}, \ \text{RowNorm}(\mathbf{0}) = \mathbf{0}. \tag{9}$$

where the multiplication aggregates the assignment weights of each sample over anchors belonging to the same anchor group. Thus, $H^{(s,g)}$ provides a soft membership representation over the anchor groups at granularity g. The sample relation induced by the transferred memberships is defined as

$$C^{(s,g)} = H^{(s,g)} \left(H^{(s,g)}\right)^{\top} \in \mathbb{R}^{N \times N}. \tag{10}$$

For the global source, the relation is normalized to obtain the global multi-granularity relation:

$$D^{(*,g)} = \text{diag}\left(C^{(*,g)} \mathbf{1}_N\right), \quad R^{(*,g)} = \left(D^{(*,g)}\right)^{-1/2} C^{(*,g)} \left(D^{(*,g)}\right)^{-1/2}. \tag{11}$$

The normalized relation can be represented in low-dimensional form as

$$Q_g = \left(D^{(*,g)}\right)^{-1/2} H^{(*,g)}, \quad R^{(*,g)} = Q_g Q_g^{\top}. \tag{12}$$

For fixed anchor groups, transfer is invariant to assignment redistributions that preserve each sample's total mass in each group. Excessive grouping can remove distinctions between samples. Appendix A.1 quantifies the effects of assignment and partition changes on the global relations.

Local Relation Refinement. For each view v, we construct a mutual-kNN graph from the observed samples and use it to refine the sample relations $C^{(v,g)}$. For an integer $k \geq 1$, let $k_v = \min\{k, n_v - 1\}$ and let $\mathcal{N}_v(i) \subseteq \Omega_v \setminus \{i\}$ contain the k_v nearest neighbors of sample i in the standardized feature space. The bandwidth is computed from the multiset of neighbor distances:

$$\tau_v = \max\left\{\text{median}\left\{\|\widetilde{x}_i^{(v)} - \widetilde{x}_j^{(v)}\|_2 : i \in \Omega_v,\ j \in \mathcal{N}_v(i)\right\}, \epsilon_\tau\right\}, \quad \epsilon_\tau > 0. \tag{13}$$

The local geometric affinity matrix $A^{(v)} \in \mathbb{R}^{N\times N}$ is initialized to zero. For $i, j \in \Omega_v$, set

$$A_{ij}^{(v)} = \exp\left(-\frac{\|\widetilde{x}_i^{(v)} - \widetilde{x}_j^{(v)}\|_2^2}{2\tau_v^2}\right), \quad (i = j) \ \vee \ \big(j \in \mathcal{N}_v(i) \ \wedge \ i \in \mathcal{N}_v(j)\big). \tag{14}$$

For each granularity g, we multiply $C^{(v,g)}$ element-wise by $A^{(v)}$ and normalize the result:

$$\widetilde{R}^{(v,g)} = C^{(v,g)} \odot A^{(v)}, \quad D^{(v,g)} = \text{diag}\Big(\widetilde{R}^{(v,g)}\mathbf{1}_N\Big), \tag{15}$$

$$R^{(v,g)} = \text{InvSqrt}_0\Big(D^{(v,g)}\Big)\,\widetilde{R}^{(v,g)}\,\text{InvSqrt}_0\Big(D^{(v,g)}\Big), \tag{16}$$

where $\odot$ denotes element-wise multiplication. For a nonnegative diagonal matrix D, $\text{InvSqrt}_0(D)$ replaces each positive diagonal entry D_{ii} with $1/\sqrt{D_{ii}}$ and keeps zero entries unchanged. Observed samples always have strictly positive degrees because their self-loop weights satisfy $\widetilde{R}_{ii}^{(v,g)} = \|H_{i,:}^{(v,g)}\|_2^2 > 0$. Rows and columns corresponding to unobserved samples remain zero.

3.3 Adaptive Low-Rank Consensus Learning

The preceding relation transfer produces VG local relations $\{R^{(v,g)}\}$ and G global relations $\{R^{(*,g)}\}$. Their quality may vary across views and granularities because they are constructed from different information and structural resolutions. We therefore learn a shared consensus relation while adaptively weighting these relation sources according to their agreement with the consensus.

Consensus Objective. Let $F \in \mathbb{R}^{N\times K}$ denote the consensus embedding. The matrix $FF^\top$ represents a positive semidefinite consensus relation with rank at most K. For each source $s \in \mathcal{S}$ and granularity g, let $P_{ij}^{(v)} = \mathbb{I}[A_{ij}^{(v)} > 0]$ and $P_{ij}^{(*)} = 1$ indicate the available sample pairs involved in consensus learning. We measure the discrepancy between $FF^\top$ and each relation source as

$$d_{s,g} = \|R^{(s,g)}\|_F^2 + \epsilon_r, \quad e_{s,g}(F) = \frac{\|P^{(s)} \odot (FF^\top - R^{(s,g)})\|_F^2}{d_{s,g}} + \epsilon_w, \tag{17}$$

where $\epsilon_r, \epsilon_w > 0$. The local masks restrict fitting to observed graph edges, and $d_{s,g}$ normalizes each error by the relation norm. We optimize the representation F by minimizing the relation discrepancy.

$$\begin{aligned} \min_{F,\alpha,\beta} \quad & \mathcal{J}(F,\alpha,\beta) = \sum_{v=1}^{V}\sum_{g=1}^{G} \alpha_{v,g}^{\gamma} e_{v,g}(F) + \lambda \sum_{g=1}^{G} \beta_g^{\gamma} e_{*,g}(F), \\ & \text{s.t.} \quad \alpha_{v,g} \geq 0, \quad \sum_{v,g} \alpha_{v,g} = 1, \quad \beta_g \geq 0, \quad \sum_{g} \beta_g = 1. \end{aligned} \tag{18}$$

The weights $\alpha_{v,g}$ and β_g control local and global contributions, respectively; $\lambda > 0$ balances the two terms and $\gamma > 1$ controls the concentration of relation weights.

Alternating Optimization. Starting from an initial representation F^0, we alternately update the adaptive weights and the consensus representation. With F^t fixed, the weight variables are optimized under the simplex constraints, resulting in the following closed-form updates:

$$\alpha_{v,g}^t = \frac{e_{v,g}(F^t)^{-p}}{\sum_{u=1}^{V}\sum_{h=1}^{G} e_{u,h}(F^t)^{-p}}, \quad \beta_g^t = \frac{e_{*,g}(F^t)^{-p}}{\sum_{h=1}^{G} e_{*,h}(F^t)^{-p}}, \quad p = \frac{1}{\gamma - 1}. \tag{19}$$

These updates emphasize relations with smaller normalized fitting residuals, while $\epsilon_w > 0$ keeps the weights well defined. The weights quantify the agreement between each relation source and

the rank-K consensus embedding. With the weights fixed, write $\mathcal{J}_t(F) = \mathcal{J}(F, \alpha^t, \beta^t)$ and $G_t = \nabla_F \mathcal{J}_t(F^t)$. We then update the consensus embedding F using Adam:

$$m_t = b_1 m_{t-1} + (1-b_1)G_t, \quad v_t = b_2 v_{t-1} + (1-b_2)(G_t \odot G_t),$$
$$\widehat{m}_t = \frac{m_t}{1-b_1^{t+1}}, \quad \widehat{v}_t = \frac{v_t}{1-b_2^{t+1}}, \quad F^{t+1} = F^t - \eta \frac{\widehat{m}_t}{\sqrt{\widehat{v}_t} + \epsilon_a}. \tag{20}$$

Here, $m_{-1} = v_{-1} = 0$, $b_1, b_2 \in (0,1)$ are the moment decay rates, $\epsilon_a > 0$, and $\eta > 0$ is a constant learning rate. The square root and division are element-wise, and the moment estimates are retained. Local residuals are evaluated on the support of $P^{(v)}$. For global relations, $R^{(*,g)} = Q_g Q_g^\top$ gives

$$\|FF^\top - R^{(*,g)}\|_F^2 = \|F^\top F\|_F^2 + \|Q_g^\top Q_g\|_F^2 - 2\|F^\top Q_g\|_F^2, \tag{21}$$
$$(FF^\top - R^{(*,g)})F = F(F^\top F) - Q_g(Q_g^\top F). \tag{22}$$

The normalization constant is computed as $d_{*,g} = \|Q_g^\top Q_g\|_F^2 + \epsilon_r$. These identities allow objective and gradient evaluation without constructing dense $N \times N$ relation matrices.

Theoretical Properties. On a fixed anchor set and granularity schedule, Appendix A analyzes how assignment, partition, degree, and effective-weight changes perturb the global relations. It identifies the target of the global fitting term and bounds changes in its leading eigenspace under a positive reference eigengap. Exact weight updates, matrix-free gradients, and existence of a global minimizer are also established. The clusters are obtained by applying K-means to the rows of F.

4 Experiments

In our experiments, we evaluate our model on five incomplete multi-view datasets, including Mfeat, MNIST-10k, Scene-15, 100Leaves, and MSRC-v1. To verify the effectiveness of the model, nine state-of-the-art methods are adopted for comparison, including **DIMVC-HIA** (Du et al. (2026)), **DGIMVCM** (Zhang et al. (2026)), **DLPMAC** (Zhao et al. (2025a)), **FreeCSL** (Dai et al. (2025)), **RPCIC** (Yuan et al. (2024)), **ICMVC** (Chao et al. (2024)), **DVSAI** (Yu et al. (2024)), **ProImp** (Li et al. (2023)), and **DIMVC** (Xu et al. (2022)). For each dataset, missing views are randomly generated under different missing rates, and all methods are evaluated using the same incomplete-view settings. More detailed of the datasets and experimental results are provided in the Appendix.

4.1 Experimental Results

Table 1 presents the clustering results of AMGC and the compared methods under incomplete-view settings. Three widely used clustering metrics, including Accuracy (ACC), Normalized Mutual Information (NMI), and Adjusted Rand Index (ARI), are adopted for evaluation. The reported results are averaged over five independent runs. All methods are evaluated under the same incomplete-view setting. From the statistical comparisons, several observations can be drawn:

1) AMGC achieves the best overall clustering performance on most benchmark datasets and incomplete-view settings. The results demonstrate that the proposed model can effectively utilize available multi-view information and learn reliable clustering structures under incomplete observations. The consistent improvements across datasets verify its effectiveness.

2) With the increase of missing rates, most comparison methods suffer from performance degradation. AMGC maintains stable results under missing-view settings, indicating that the proposed framework can better preserve useful information when view observations become incomplete.

3) Compared with existing IMVC methods, AMGC obtains consistent improvements on several benchmark datasets. This suggests that exploring structural information from different views and resolutions is beneficial for discovering accurate cluster structures in incomplete multi-view data.

Overall, the experimental results verify the effectiveness of AMGC for incomplete multi-view clustering and demonstrate its ability to learn reliable clustering structures from partially observed views.

Table 1: Clustering performance (%) under different missing-view rates on five benchmark datasets. Best results for each dataset, missing-view rate, and metric are highlighted in bold.

Data	Method	0.1			0.3			0.5			0.7		
		ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI
Mfeat	AMGC	**95.45**	**90.20**	**90.24**	**92.60**	**85.67**	**84.47**	**91.35**	**83.25**	**81.89**	**91.05**	**81.63**	**81.38**
	DIMVC-HIA	92.85	86.62	88.54	88.35	82.79	83.31	82.77	79.21	81.03	80.15	77.22	78.99
	DGIMVCM	90.80	85.46	87.35	86.14	80.39	81.26	82.09	78.43	79.43	81.95	78.03	78.36
	DLPMAC	75.31	74.86	66.49	66.15	64.98	50.06	59.51	59.74	38.20	54.50	55.47	31.25
	FreeCSL	90.15	85.81	83.75	79.55	81.17	71.09	81.41	79.85	71.55	82.81	80.83	74.32
	RPCIC	78.95	75.07	69.53	66.35	62.46	48.67	60.80	54.82	34.74	53.35	49.16	22.41
	DIMVC	80.24	79.27	73.26	78.21	75.74	68.02	80.64	79.54	75.48	78.64	76.73	70.21
	ICMVC	79.16	72.67	62.55	73.42	68.47	56.37	68.22	61.36	48.43	62.45	55.66	41.56
	DVSAI	91.17	86.04	84.98	87.15	82.34	81.38	84.83	78.74	76.82	81.05	75.56	73.22
	ProImp	81.57	78.01	70.08	76.71	71.85	65.42	68.86	66.87	58.38	60.07	61.95	46.77
MNIST-10k	AMGC	**91.99**	**80.96**	**81.69**	**90.15**	**78.89**	**79.78**	**89.96**	**78.68**	**79.57**	**86.02**	**71.91**	**72.36**
	DIMVC-HIA	80.12	65.79	62.82	81.43	66.85	64.40	83.70	69.95	67.96	84.34	69.90	69.04
	DGIMVCM	84.88	75.71	76.91	82.33	72.49	74.60	80.12	69.43	70.89	75.75	63.84	65.03
	DLPMAC	79.86	76.71	76.66	79.93	76.89	76.57	60.74	56.08	36.24	58.39	56.34	34.96
	FreeCSL	90.30	79.57	80.04	74.11	71.91	63.85	72.32	70.48	62.13	68.42	68.24	59.70
	RPCIC	85.92	73.41	72.43	74.63	65.80	61.03	63.83	53.28	43.24	35.26	31.21	18.61
	DIMVC	84.36	72.15	71.90	88.93	76.43	74.50	80.84	70.81	68.04	39.23	39.02	23.94
	ICMVC	82.73	68.27	66.74	87.79	76.15	75.22	85.73	76.02	75.48	82.98	67.80	68.27
	DVSAI	85.33	71.85	71.16	83.44	68.78	67.89	81.10	65.53	63.93	71.61	57.36	54.43
	ProImp	87.98	76.07	75.75	71.33	63.07	54.83	65.40	57.84	40.77	61.04	54.08	39.51
Scene-15	AMGC	**51.62**	**45.06**	**32.01**	**46.55**	**43.25**	**30.76**	**41.75**	**38.78**	**28.33**	**40.56**	**36.83**	**25.73**
	DIMVC-HIA	47.23	43.07	30.96	44.07	42.81	29.28	40.63	36.27	27.99	37.21	34.91	23.13
	DGIMVCM	43.56	41.85	29.68	41.43	38.79	26.16	35.64	34.92	23.08	32.31	32.09	19.55
	DLPMAC	30.48	29.16	14.79	28.27	27.75	13.22	24.68	24.56	11.09	22.72	23.04	9.75
	FreeCSL	28.49	32.04	15.48	29.99	29.45	15.88	27.65	26.57	13.30	24.66	23.56	10.98
	RPCIC	38.39	36.32	21.23	39.69	36.91	21.65	30.12	31.19	15.11	19.51	22.22	8.07
	DIMVC	46.53	42.38	30.51	43.84	39.66	24.54	39.53	37.82	23.31	38.36	35.49	19.43
	ICMVC	42.49	40.54	25.75	41.27	38.72	24.40	38.70	35.61	22.04	38.33	32.91	18.05
	DVSAI	40.09	39.68	23.19	38.51	36.16	20.38	35.63	32.27	18.62	35.18	31.14	18.45
	ProImp	36.84	38.87	22.07	31.97	32.61	17.45	29.73	29.17	13.95	26.68	27.06	11.42
100Leaves	AMGC	**88.25**	**94.31**	**83.46**	**82.44**	**90.43**	**73.68**	**77.78**	**87.37**	**66.28**	**71.19**	**83.86**	**58.51**
	DIMVC-HIA	87.51	90.66	81.72	80.31	86.82	72.35	76.40	82.21	63.86	69.19	79.12	55.64
	DGIMVCM	81.44	85.81	74.28	76.19	80.42	69.08	75.88	79.38	64.69	69.25	75.81	57.76
	DLPMAC	79.31	89.45	71.49	58.75	77.67	43.22	41.19	66.19	22.63	30.44	59.90	13.76
	FreeCSL	74.31	87.19	64.54	61.31	80.38	49.24	46.62	71.21	32.26	34.94	63.38	19.80
	RPCIC	76.62	88.06	67.13	61.75	79.88	46.79	48.44	70.83	27.80	38.19	65.68	18.57
	DIMVC	83.31	90.65	81.17	71.59	88.63	68.33	27.34	65.17	21.74	11.59	45.27	15.21
	ICMVC	46.21	70.93	25.22	37.15	60.90	19.87	32.19	54.33	14.56	26.69	47.61	11.41
	DVSAI	75.31	89.48	64.74	65.75	82.46	50.82	56.44	76.24	39.48	46.38	69.74	27.36
	ProImp	61.88	87.03	57.91	40.34	74.26	33.43	33.88	66.60	23.16	26.38	60.60	13.84
MSRC_v1	AMGC	**90.48**	**81.53**	**79.51**	**84.29**	**70.67**	**67.86**	**76.19**	**65.84**	**60.11**	**72.38**	**63.74**	**53.48**
	DIMVC-HIA	71.62	59.89	46.17	66.48	55.66	41.17	61.33	46.51	36.08	57.52	42.46	32.34
	DGIMVCM	82.10	76.82	73.36	76.00	68.96	64.81	73.81	63.34	59.82	65.90	59.02	51.56
	DLPMAC	69.52	71.82	63.51	66.67	63.99	50.95	59.05	51.71	36.32	55.71	50.58	33.56
	FreeCSL	79.52	75.59	67.29	73.33	66.46	57.28	70.86	60.83	54.44	65.24	62.51	50.12
	RPCIC	73.81	63.29	55.93	48.57	36.36	28.12	32.86	21.62	21.23	28.86	15.89	17.51
	DIMVC	66.00	56.23	44.86	62.38	51.56	38.35	63.14	53.52	42.51	61.24	50.01	36.88
	ICMVC	75.43	65.15	58.05	36.00	21.15	20.64	30.67	16.62	16.85	27.81	11.95	13.69
	DVSAI	74.76	65.50	57.05	71.90	61.14	53.26	70.48	53.05	42.16	66.67	53.47	43.31
	ProImp	73.33	66.05	56.81	72.95	63.64	53.61	63.05	53.45	37.04	59.62	47.99	28.35

4.2 Ablation Study

To evaluate the effectiveness of different components in AMGC, we conduct an ablation study under the missing rate of 0.5. The results are reported in Table 2, and several observations can be drawn:

1) According to the results, removing any component leads to performance degradation on most datasets, demonstrating that all components contribute to the clustering performance of AMGC. Among them, removing the LG module causes an obvious drop, indicating that preserving local geometric information is important for maintaining sample relations under incomplete observations.

2) Removing the SAR module also results in noticeable performance degradation on several datasets. This demonstrates that constructing a shared relational reference is beneficial for establishing a unified reference space across incomplete views and providing consistent structural information for subsequent relation learning.

Table 2: Ablation study of different components of AMGC under the missing rate of 0.5 (%).

SAR	MGR	LG	ARC	Mfeat	MNIST-10k	Scene-15	100Leaves	MSRC-v1
	✓	✓	✓	68.65±2.40	87.98±0.24	40.62±0.34	71.96±1.50	74.83±2.34
✓		✓	✓	80.16±0.31	81.82±0.09	39.93±1.67	72.37±0.88	69.71±2.62
✓	✓		✓	83.62±0.54	66.84±3.28	25.98±0.58	44.38±1.73	45.87±3.44
✓	✓	✓		88.78±0.42	87.94±0.08	40.25±0.17	75.91±0.36	74.67±1.52
✓	✓	✓	✓	**91.35±0.52**	**89.96±0.04**	**41.75±0.32**	**77.78±0.45**	**76.19±1.44**

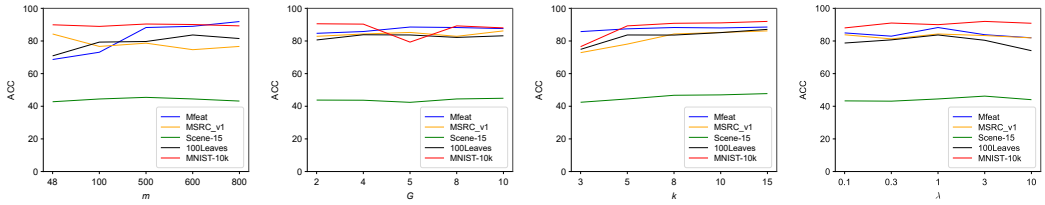

(a) Anchor number m (b) Granularity number G (c) Neighborhood size k (d) Balance coefficient λ

Figure 2: Hyperparameter sensitivity of AMGC on five benchmark datasets at a missing rate of 0.3. Each curve represents one dataset, and ACC is reported as a percentage.

3) The variants without MGR or ARC also show performance decreases compared with the complete model. These results verify that modeling structures across different granularities and adaptively integrating multiple relational sources are helpful for discovering more accurate cluster structures.

4.3 Hyperparameter Analysis

We investigate the sensitivity of AMGC to four hyperparameters, including the anchor number m, granularity number G, neighborhood size k, and global relation coefficient λ. The anchor number is capped at the number of samples N, i.e., $m \leftarrow \min(m, N)$. Experiments are conducted on five benchmark datasets with a missing rate of 0.3, and the results are shown in Fig. 2.

1) Anchor number m: The number of anchors affects the quality of shared reference construction. Insufficient anchors may limit structural modeling, while excessive anchors provide limited benefits.

2) Granularity number G: The granularity number controls the ability to capture multi-scale structures. An appropriate number of granularities achieves a balance between fine and coarse structures.

3) Neighborhood size k: The neighborhood size k determines the range of local geometric information captured by the mutual k-nearest-neighbor graphs. A small k may fail to capture sufficient local structures, while an excessively large k may introduce less relevant neighbors.

4) Global relation coefficient λ: The coefficient λ balances local and global relations, and stable performance under different settings demonstrates the robustness of adaptive consensus learning.

Overall, AMGC maintains robust clustering performance across a wide range of hyperparameters.

5 Conclusion

Incomplete observations introduce inconsistent relational references across views and unstable cluster structures across different granularities, which limit the effectiveness of multi-granularity modeling for incomplete multi-view clustering. To address these issues, we propose AMGC, a novel anchor-induced multi-granularity consensus framework. AMGC introduces shared anchors to construct a unified relational reference across incomplete views and develops an anchor-induced relation transfer mechanism to capture fine-to-coarse cluster structures with sparse local geometric guidance. Furthermore, an adaptive low-rank consensus learning strategy is designed to integrate different view-granularity relations according to their consistency with the learned consensus relation. Extensive experiments on multiple benchmark incomplete multi-view datasets demonstrate the effectiveness of the proposed model over existing IMVC methods.

AI USE STATEMENT

Generative AI tools were used solely for language refinement. The authors reviewed and revised the manuscript to ensure the accuracy and quality of the final version.

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*, volume 38, pp. 11221–11229, 2024.

Guoqing Chao, Kaixin Xu, Xijiong Xie, and Yongyong Chen. Global graph propagation with hierarchical information transfer for incomplete contrastive multi-view clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 15713–15721, 2025a.

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*, volume 267, pp. 7417–7429, 2025b.

Feiyang Chen, Ruiqiang Guo, and Zhibin Gu. Multi-scale diffusion-guided graph learning with power-smoothing random walk contrast for multi-view clustering. In *International Conference on Learning Representations*, 2026.

Yongyong Chen, Xiaojia Zhao, Zheng Zhang, Youfa Liu, Jingyong Su, and Yicong Zhou. Tensor learning meets dynamic anchor learning: from complete to incomplete multiview clustering. *IEEE Transactions on Neural Networks and Learning Systems*, 35(11):15332–15345, 2024.

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.

Fengyang Ding, Hui Huang, Min Yu, Wei Jiang, Zhen Xu, Jun Yin, and Nan Zhang. Filtering and denoising: graph completion-based incomplete multi-view clustering. *Information Fusion*, 126: 103654, 2026.

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*, volume 40, pp. 20941–20949, 2026.

Lele Fu, Pengfei Lin, Athanasios V. Vasilakos, and Shiping Wang. An overview of recent multi-view clustering. *Neurocomputing*, 402:148–161, 2020.

Teofilo F. Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science*, 38:293–306, 1985.

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.

Wei Guo, Zhe Wang, and Wei Shao. Structure anchor graph learning for multi-view clustering. *Pattern Recognition*, 170:111880, 2026.

Wen-Jue He, Zheng Zhang, and Xiaofeng Zhu. Dual-correlation-guided anchor learning for scalable incomplete multi-view clustering. *IEEE Transactions on Neural Networks and Learning Systems*, 36(9):17336–17349, 2025.

Menglei Hu and Songcan Chen. One-pass incomplete multi-view clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 3838–3845, 2019.

Enze Ji, Xuan Chen, Zhikui Chen, and Chuan Xiao. Incomplete multi-view clustering via imbalanced anchor-aligned graph learning. *Pattern Recognition*, 180:114290, 2026.

Jintian Ji and Songhe Feng. Anchors crash tensor: efficient and scalable tensorial multi-view subspace clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47:2660–2675, 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.

Haobin Li, Yunfan Li, Mouxing Yang, Peng Hu, Dezhong Peng, and Xi Peng. Incomplete multi-view clustering via prototype-based imputation. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pp. 3911–3919, 2023.

Xiang-Long Li, Man-Sheng Chen, Chang-Dong Wang, and Jian-Huang Lai. Refining graph structure for incomplete multi-view clustering. *IEEE Transactions on Neural Networks and Learning Systems*, 35:2300–2313, 2024.

Zhenglai Li, Chang Tang, Xiao Zheng, Xinwang Liu, Wei Zhang, and En Zhu. High-order correlation preserved incomplete multi-view subspace clustering. *IEEE Transactions on Image Processing*, 31:2067–2080, 2022.

Yijie Lin, Yuanbiao Gou, Zitao Liu, Boyun Li, Jiancheng Lv, and Xi Peng. Completer: Incomplete multi-view clustering via contrastive prediction. In *2021 IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, pp. 11169–11178, 2021.

Yijie Lin, Yuanbiao Gou, Xiaotian Liu, Jinfeng Bai, Jiancheng Lv, and Xi Peng. Dual contrastive prediction for incomplete multi-view representation learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45:4447–4461, 2023.

Jiyuan Liu, Xinwang Liu, Siqi Wang, Xinhang Wan, Dongsheng Li, Kai Lu, and Kunlun He. Communication-efficient federated multi-view clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 48:17–32, 2026.

Zhen Long, Qiyuan Wang, Yazhou Ren, Yipeng Liu, and Ce Zhu. TLRLF4MVC: tensor low-rank and low-frequency for scalable multi-view clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47:6900–6911, 2025.

Jingyu Pu, Chenhang Cui, Xinyue Chen, Yazhou Ren, Xiaorong Pu, Zhifeng Hao, Philip S. Yu, and Lifang He. Adaptive feature imputation with latent graph for deep incomplete multi-view clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 14633–14641, 2024.

Yalan Qin, Guorui Feng, and Xinpeng Zhang. Scalable one-pass incomplete multi-view clustering by aligning anchors. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 20042–20050, 2025a.

Yalan Qin, Xinpeng Zhang, Shui Yu, and Guorui Feng. A survey on representation learning for multi-view data. *Neural Networks*, 181:106842, 2025b.

Qianqian Wang, Haiming Xu, Meiling Liu, Wei Feng, and Xiangdong Zhang. Fair incomplete multi-view clustering via distribution alignment. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, pp. 6379–6387, 2025a.

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

Suixue Wang, Shilin Zhang, Qingchen Zhang, Peng Li, and Weiliang Huo. MASTER: a multi-granularity invariant structure clustering scheme for multi-view clustering. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, pp. 6415–6423, 2025b.

Youqing Wang, Yu Cao, Jinlu Wang, Xiang Xu, Jiapu Wang, Tengfei Liu, Junbin Gao, and Jipeng Guo. Aligning collaborative view recovery and tensorial subspace learning via latent representation for incomplete multi-view clustering. In *Proceedings of the International Conference on Learning Representations*, pp. 79726–79740, 2026b.

Xiaohui Wei, Yuting Chen, Feiping Nie, Haibo Liu, Qiya Song, and Lin Xiao. View-adaptive multi-granularity anchor learning for multi-view clustering. *IEEE Transactions on Image Processing*, 35:3113–3126, 2026.

Jie Wen, Zheng Zhang, Yong Xu, Bob Zhang, Lunke Fei, and Hong Liu. Unified embedding alignment with missing views inferring for incomplete multi-view clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 5393–5400, 2019.

Jie Wen, Yong Xu, and Hong Liu. Incomplete multiview spectral clustering with adaptive graph learning. *IEEE Transactions on Cybernetics*, 50:1418–1429, 2020.

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 *Forty-first international conference on machine learning*, 2024.

Jie Xu, Chao Li, Yazhou Ren, Liang Peng, Yujie Mo, Xiaoshuang Shi, and Xiaofeng Zhu. Deep incomplete multi-view clustering via mining cluster complementarity. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pp. 8761–8769, 2022.

Ben Yang, Xuetao Zhang, Zhiyuan Xue, Feiping Nie, and Badong Chen. Fast multi-view discrete clustering via spectral embedding fusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 48:5095–5112, 2026.

Jie Yang, Wei Chen, Feng Liu, Peng Zhou, Zhongli Wang, Xinyan Liang, and Bingbing Jiang. Multi-view clustering via multi-granularity ensemble. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, pp. 6794–6802, 2025.

Shengju Yu, Siwei Wang, Pei Zhang, Miao Wang, Ziming Wang, Zhe Liu, Liming Fang, En Zhu, and Xinwang Liu. DVSAI: Diverse view-shared anchors based incomplete multi-view clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 16568–16577, 2024.

Shengju Yu, Zhibin Dong, Siwei Wang, Pei Zhang, Yi Zhang, Xinwang Liu, Naiyang Guan, Tiejun Li, and Yiu ming Cheung. Simple yet effective incomplete multi-view clustering: similarity-level imputation and intra-view hybrid-group prototype construction. In *Proceedings of the International Conference on Learning Representations*, pp. 6971–6997, 2025.

Honglin Yuan, Shiyun Lai, Xingfeng Li, Jian Dai, Yuan Sun, and Zhenwen Ren. Robust prototype completion for incomplete multi-view clustering. In *Proceedings of the 32nd ACM international conference on multimedia*, pp. 10402–10411, 2024.

Pengxin Zeng, Mouxing Yang, Yiding Lu, Changqing Zhang, Peng Hu, and Xi Peng. Semantic invariant multi-view clustering with fully incomplete information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(4):2139–2150, 2024.

Chao Zhang, Huaxiong Li, Caihua Chen, Xiuyi Jia, and Chunlin Chen. Low-rank tensor regularized views recovery for incomplete multiview clustering. *IEEE Transactions on Neural Networks and Learning Systems*, 35(7):9312–9324, 2024.

Zhenghao Zhang, Jun Xie, Xingchen Chen, Tao Yu, Hongzhu Yi, Kaixin Xu, Yuanxiang Wang, Tianyu Zong, Xinming Wang, Jiahuan Chen, et al. Dynamic deep graph learning for incomplete multi-view clustering with masked graph reconstruction loss. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 28600–28608, 2026.

Liang Zhao, Shubin Ma, Bo Xu, and Qingchen Zhang. Dual-learning based penalized multi-align clustering for multi-view incomplete and disorderly data. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pp. 1318–1326, 2025a.

Shuping Zhao, Jie Wen, Lunke Fei, and Bob Zhang. Tensorized incomplete multi-view clustering with intrinsic graph completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 11327–11335, 2023.

Shuping Zhao, Lunke Fei, Qi Lai, Jie Wen, Jinrong Cui, and Tingting Chai. High-confident local structure guided consensus graph learning for incomplete multi-view clustering. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, pp. 7119–7127, 2025b.

Shuping Zhao, Yulong Chen, Jie Wen, Lunke Fei, Jinrong Cui, and Tingting Chai. Confident block diagonal structure-aware invariable graph completion for incomplete multi-view clustering. In *International Conference on Learning Representations*, pp. 99860–99870, 2026.

A Theoretical Analysis and Optimization Properties

Sections A.1 and A.2 analyze anchor-induced relation transfer for the global source. Sections A.3–A.5 analyze optimization with fixed input relations and masks. All $R^{(s,g)}$ and $P^{(s)}$ are symmetric, $P^{(s)}$ is binary, and $P^{(*)} = \mathbf{1}_N \mathbf{1}_N^\top$. We assume $N, K, V, G \geq 1$, $\lambda > 0$, $\gamma > 1$, and $\epsilon_r, \epsilon_w > 0$, with $F \in \mathbb{R}^{N \times K}$ unconstrained. Let $\mathcal{W}$ be the product of the two weight simplexes in Eq. 18.

A.1 Perturbation of Anchor-Induced Relation Transfer

We compare two global constructions on the same samples and selected anchor identities, with the same granularity sizes. Write $Z = Z^{(*)}$ and $Y_g = Y^{(*,g)}$, and use hats for the second construction. The assignments are nonnegative with positive row sums; both $Y_g, \widehat{Y}_g \in \{0,1\}^{m \times c_g}$ have one nonzero per row. AMGC uses unit-row-sum assignments. To allow a more general perturbation setting, we consider perturbed constructions with positive but nonunit row sums. Let $\mathcal{N}$ divide each positive row by its sum, with $\mathcal{N}(0) = 0$. Following Eq. 9, define

$$B_g = ZY_g, \quad s_g = B_g \mathbf{1}_{c_g}, \quad S_g = \operatorname{diag}(s_g), \quad H_g = \mathcal{N}(B_g) = S_g^{-1} B_g, \quad C_g = H_g H_g^\top. \tag{23}$$

Set $D_g = \operatorname{diag}(C_g \mathbf{1}_N)$ and $R_g = D_g^{-1/2} C_g D_g^{-1/2}$; hatted quantities are defined likewise. Partition labels may be aligned by a column permutation, which leaves the corresponding relation unchanged. Since H_g is nonnegative with unit row sums, $(D_g)_{ii} \geq \|(H_g)_{i,:}\|_2^2 \geq 1/c_g$.

Proposition. Let $E = \widehat{Z} - Z$ and define

$$\begin{aligned} t_g &= \|EY_g\|_F + \|\widehat{Z}(\widehat{Y}_g - Y_g)\|_F, \\ \underline{s}_g &= \min_i \{s_{g,i}, \widehat{s}_{g,i}\} > 0, \quad u_g = \frac{t_g + \|\widehat{s}_g - s_g\|_2}{\underline{s}_g}, \\ b_g &= (2\|H_g\|_2 + u_g) u_g, \quad \underline{d}_g = \min_i (D_g)_{ii}, \\ \theta_g &= \max_i \frac{|(\widehat{D}_g)_{ii} - (D_g)_{ii}|}{(D_g)_{ii}}. \end{aligned} \tag{24}$$

Then $\|\widehat{H}_g - H_g\|_F \leq u_g$ and $\|\widehat{C}_g - C_g\|_F \leq b_g$. If $\theta_g < 1$, then

$$\|\widehat{R}_g - R_g\|_F \leq \frac{b_g/\underline{d}_g + \theta_g \|R_g\|_F}{1 - \theta_g} =: \epsilon_g. \tag{25}$$

For AMGC's unit-row-sum assignments, $u_g = t_g$. If $\widehat{Y}_g = Y_g$ and $EY_g = 0$, then $\widehat{R}_g = R_g$.

Proof. Before row normalization, the exact perturbation is

$$\Delta B_g := \widehat{B}_g - B_g = EY_g + \widehat{Z}(\widehat{Y}_g - Y_g), \tag{26}$$

so $\|\Delta B_g\|_F \leq t_g$. Since $B_g = S_g H_g$, row normalization gives

$$\widehat{H}_g - H_g = \widehat{S}_g^{-1} \big[\Delta B_g - (\widehat{S}_g - S_g) H_g\big]. \tag{27}$$

Each row of H_g has Euclidean norm at most one. Hence $\|(\widehat{S}_g - S_g) H_g\|_F \leq \|\widehat{s}_g - s_g\|_2$; together with $\|\widehat{S}_g^{-1}\|_2 \leq 1/\underline{s}_g$, this proves the bound by u_g. Moreover, $\widehat{s}_g - s_g = \Delta B_g \mathbf{1}_{c_g}$ implies $u_g \leq (1 + \sqrt{c_g}) t_g / \underline{s}_g$. Expanding the Gram matrices now gives $\|\widehat{C}_g - C_g\|_F \leq (2\|H_g\|_2 + u_g) u_g$.

For degree normalization, set $L_g = \widehat{D}_g^{-1/2} D_g^{1/2}$. Then

$$\widehat{R}_g - R_g = \widehat{D}_g^{-1/2} (\widehat{C}_g - C_g) \widehat{D}_g^{-1/2} + L_g R_g L_g - R_g. \tag{28}$$

For $\theta_g < 1$, the first term has Frobenius norm at most $b_g / [(1 - \theta_g) \underline{d}_g]$. Also, $\|L_g\|_2 \leq (1 - \theta_g)^{-1/2}$ and $\|L_g - I\|_2 \leq (1 - \theta_g)^{-1/2} - 1$, giving

$$\|L_g R_g L_g - R_g\|_F \leq \|L_g - I\|_2 (\|L_g\|_2 + 1) \|R_g\|_F \leq \frac{\theta_g}{1 - \theta_g} \|R_g\|_F. \tag{29}$$

Combining the two bounds proves Eq. 25. For the AMGC construction, $s_g = ZY_g\mathbf{1}_{c_g} = Z\mathbf{1}_m = \mathbf{1}_N$, and similarly $\widehat{s}_g = \mathbf{1}_N$, so $u_g = t_g$. Finally, $\widehat{Y}_g = Y_g$ and $EY_g = 0$ imply $\widehat{B}_g = B_g$; applying the same row and degree normalizations therefore gives $\widehat{R}_g = R_g$. □

The invariance result characterizes a special case where assignment redistribution occurs within fixed anchor partitions while preserving group-level assignment masses. More generally, the perturbation bound in Eq. 25 quantifies the effects of both assignment and partition changes through t_g and θ_g. At the extreme $c_g = 1$, $R_g = \mathbf{1}_N\mathbf{1}_N^\top/N$ is invariant to positive-row-sum Z but retains no distinctions between samples.

A.2 Global Fitting Target and Spectral Stability

For fixed global weights, we analyze the target of the global fitting term by rewriting it as a weighted combination of multi-granularity relations. Let $\omega_g = \beta_g^\gamma/d_{*,g}$, $\omega = \sum_g \omega_g > 0$, $\pi_g = \omega_g/\omega$, and $R_\pi = \sum_g \pi_g R_g$ denote the induced weights and consensus relation. Completing the square gives

$$\sum_g \beta_g^\gamma e_{*,g}(F) = \omega\|FF^\top - R_\pi\|_F^2 + C_\beta, \quad C_\beta = \sum_g \omega_g\|R_g\|_F^2 - \omega\|R_\pi\|_F^2 + \epsilon_w\sum_g \beta_g^\gamma. \tag{30}$$

where C_β is independent of F. This identifies the effective target of the global fitting term without altering the objective. To further characterize this target, let $Q_g = D_g^{-1/2}H_g$ as in Eq. 12, which gives $R_g = Q_gQ_g^\top$. Then, $\mathcal{Q}_\pi = [\sqrt{\pi_1}Q_1, \ldots, \sqrt{\pi_G}Q_G]$ and $R_\pi = \mathcal{Q}_\pi\mathcal{Q}_\pi^\top$, with $\mathrm{range}(R_\pi) = \mathrm{range}(\mathcal{Q}_\pi)$. The range equality follows from $\ker(\mathcal{Q}_\pi\mathcal{Q}_\pi^\top) = \ker(\mathcal{Q}_\pi^\top)$. Thus, the global fitting target combines the relational information from all granularities with positive weights.

Corollary. Let $\widehat{\pi}$ be the effective weights of the second construction, and set $\widehat{R}_{\widehat{\pi}} = \sum_g \widehat{\pi}_g\widehat{R}_g$. To characterize the stability of the corresponding spectral subspaces, suppose that $\|\widehat{R}_g - R_g\|_F \le \epsilon_g$ for every g. For $1 \le K < N$, assume $\delta = \lambda_K(R_\pi) - \lambda_{K+1}(R_\pi) > 0$, with eigenvalues in decreasing order. If $U, \widehat{U}$ contain the leading K orthonormal eigenvectors of R_π and $\widehat{R}_{\widehat{\pi}}$, respectively, then

$$\|\sin\Theta(\widehat{U}, U)\|_F \le \frac{2}{\delta}\left(\sum_g \widehat{\pi}_g\epsilon_g + \sum_g |\widehat{\pi}_g - \pi_g|\,\|R_g\|_F\right), \tag{31}$$

where Θ contains the principal angles between the two subspaces.

Proof. Write $A = R_\pi$, $\widehat{A} = \widehat{R}_{\widehat{\pi}}$, and $\Delta_\pi = \widehat{A} - A$. The resulting perturbation can be decomposed as

$$\Delta_\pi = \sum_g \widehat{\pi}_g(\widehat{R}_g - R_g) + \sum_g(\widehat{\pi}_g - \pi_g)R_g. \tag{32}$$

Taking the Frobenius norm of this decomposition and applying the triangle inequality gives

$$\|\Delta_\pi\|_F \le \underbrace{\sum_g \widehat{\pi}_g\epsilon_g}_{\text{relation construction error}} + \underbrace{\sum_g |\widehat{\pi}_g - \pi_g|\,\|R_g\|_F}_{\text{adaptive weight variation}}. \tag{33}$$

Set $P = UU^\top$, $\widehat{P} = \widehat{U}\widehat{U}^\top$, and $t = \|\sin\Theta(\widehat{U}, U)\|_F$. By the definition of principal angles,

$$t^2 = K - \mathrm{tr}(P\widehat{P}) = \tfrac{1}{2}\|P - \widehat{P}\|_F^2. \tag{34}$$

Let $v_1, \ldots, v_N$ be an orthonormal eigenbasis of A, ordered by $\lambda_1 \ge \cdots \ge \lambda_N$, and set $q_i = v_i^\top\widehat{P}v_i$. Then $0 \le q_i \le 1$, $\sum_i q_i = K$, and $\sum_{i>K} q_i = t^2$. Consequently,

$$\mathrm{tr}\big(A(P - \widehat{P})\big) = \sum_{i\le K}\lambda_i(1 - q_i) - \sum_{i>K}\lambda_i q_i \ge (\lambda_K - \lambda_{K+1})t^2 = \delta t^2. \tag{35}$$

By the variational characterization of the leading eigenspace of $\widehat{A}$,

$$\mathrm{tr}(\widehat{A}\widehat{P}) = \sum_{i=1}^K \lambda_i(\widehat{A}) \ge \mathrm{tr}(\widehat{A}P). \tag{36}$$

Using $\widehat{A} = A + \Delta_\pi$ and the Frobenius Cauchy–Schwarz inequality, we obtain

$$\delta t^2 \le \mathrm{tr}\big(\Delta_\pi(\widehat{P} - P)\big) \le \|\Delta_\pi\|_F \|\widehat{P} - P\|_F = \sqrt{2}\|\Delta_\pi\|_F t. \tag{37}$$

If $t = 0$, the claim is immediate. Otherwise, dividing by δt gives $t \le \frac{\sqrt{2}}{\delta}\|\Delta_\pi\|_F$. Substituting the bound on $\|\Delta_\pi\|_F$ and using $\sqrt{2} < 2$ yields Eq. 31. □

The second term captures the effect of changes in the effective weights, including variations induced by $d_{*,g}$. Together with the previous perturbation bound, this result shows that the global fitting target remains spectrally stable when the transferred relations and their weights undergo bounded changes. This analysis focuses on the stability of the anchor-induced consensus target and its spectral structure, providing a theoretical characterization of the global relation component of AMGC.

A.3 Exact Weight Minimization

For fixed F, consider a group of q positive errors $e_1, \dots, e_q$ and weights $w_a \ge 0$ with $\sum_a w_a = 1$. Set $p = 1/(\gamma - 1)$ and $Z = \sum_a e_a^{-p}$. Hölder's inequality gives

$$\begin{aligned} 1 &= \left(\sum_{a=1}^{q} (e_a^{1/\gamma} w_a) e_a^{-1/\gamma}\right)^{\gamma} \le \left(\sum_{a=1}^{q} e_a w_a^{\gamma}\right) Z^{\gamma-1}, \\ \sum_{a=1}^{q} e_a w_a^{\gamma} &\ge Z^{1-\gamma} = Z^{-\gamma} \sum_{a=1}^{q} e_a^{-p} = \sum_{a=1}^{q} e_a \left(\frac{e_a^{-p}}{Z}\right)^{\gamma}. \end{aligned} \tag{38}$$

The bound is attained by $w_a = e_a^{-p}/Z$, giving the optimal weight solution. Strict convexity for $\gamma > 1$ ensures uniqueness. Applying this result to local and global weights yields

$$\begin{aligned} Z_\alpha(F) &= \sum_{v,g} e_{v,g}(F)^{-p}, \quad Z_\beta(F) = \sum_{g} e_{*,g}(F)^{-p}, \\ \alpha_{v,g}(F) &= \frac{e_{v,g}(F)^{-p}}{Z_\alpha(F)}, \qquad \beta_g(F) = \frac{(\lambda e_{*,g}(F))^{-p}}{\sum_h (\lambda e_{*,h}(F))^{-p}} = \frac{e_{*,g}(F)^{-p}}{Z_\beta(F)}. \end{aligned} \tag{39}$$

Since $e_{s,g}(F) \ge \epsilon_w > 0$, these weight maps are continuous and strictly positive. After exact minimization over the adaptive weights, the objective can be reduced to

$$\Phi(F) = \min_{(\alpha,\beta)\in\mathcal{W}} \mathcal{J}(F, \alpha, \beta) = \mathcal{J}(F, \alpha(F), \beta(F)) = Z_\alpha(F)^{1-\gamma} + \lambda Z_\beta(F)^{1-\gamma}. \tag{40}$$

Remark (Rank-dependent fitting difficulty). For a global relation $R_g = R^{(*,g)} = Q_g Q_g^\top \succeq 0$, let $\mu_1 \ge \cdots \ge \mu_N \ge 0$ denote its eigenvalues. Under the parameterization $F \in \mathbb{R}^{N\times K}$,

$$\min_F e_{*,g}(F) = \epsilon_w + \zeta_g, \quad \zeta_g := \frac{\sum_{j>K} \mu_j^2}{\|R_g\|_F^2 + \epsilon_r}. \tag{41}$$

Indeed, let P_F be the orthogonal projector onto the column space of F. Since $(I - P_F)FF^\top = 0$ and $\mathrm{rank}(P_F) \le K$,

$$\|FF^\top - R_g\|_F^2 \ge \mathrm{tr}(R_g^2) - \mathrm{tr}(P_F R_g^2) \ge \sum_{j>K} \mu_j^2. \tag{42}$$

The last inequality follows because the diagonal entries of P_F in an eigenbasis of R_g lie in $[0, 1]$ and their sum is bounded by K. Retaining the largest at most K eigenvalues of R_g attains this bound and admits a corresponding factorization $FF^\top$, with zero columns added when necessary.

Thus, $c_g \le K$ implies $\zeta_g = 0$, whereas $c_g > K$ does not necessarily imply $\zeta_g > 0$. These are individual-source minima and need not be attained simultaneously by a common F. Consequently, inverse-error weighting may favor relations that are easier to approximate at rank K, rather than semantically more reliable relations. This observation alone does not establish collapse of the model.

A.4 Embedding Gradient

For any source s and granularity g, define $E_{s,g}(F) = P^{(s)} \odot (FF^\top - R^{(s,g)})$. For an arbitrary perturbation $U \in \mathbb{R}^{N\times K}$, the symmetry of $P^{(s)}$ and $P^{(s)} \odot P^{(s)} = P^{(s)}$ give

$$\begin{aligned} \mathrm{D}e_{s,g}(F)[U] &= \frac{2}{d_{s,g}} \left\langle E_{s,g}(F), P^{(s)} \odot (UF^\top + FU^\top) \right\rangle_F \\ &= \frac{2}{d_{s,g}} \left\langle E_{s,g}(F), UF^\top + FU^\top \right\rangle_F \\ &= \frac{4}{d_{s,g}} \left\langle E_{s,g}(F)F, U \right\rangle_F, \quad \nabla_F e_{s,g}(F) = \frac{4}{d_{s,g}} E_{s,g}(F)F. \end{aligned} \tag{43}$$

For the global relation, let $R^{(*,g)} = Q_g Q_g^\top$ as defined in Eq. 12. The corresponding terms can be evaluated without explicitly constructing dense $N \times N$ relations:

$$\begin{aligned} \|FF^\top - Q_g Q_g^\top\|_F^2 &= \mathrm{tr}\big((FF^\top)^2\big) + \mathrm{tr}\big((Q_g Q_g^\top)^2\big) - 2\,\mathrm{tr}(FF^\top Q_g Q_g^\top) \\ &= \|F^\top F\|_F^2 + \|Q_g^\top Q_g\|_F^2 - 2\|F^\top Q_g\|_F^2, \end{aligned} \tag{44}$$

and

$$E_{*,g}(F)F = (FF^\top - Q_g Q_g^\top)F = F(F^\top F) - Q_g(Q_g^\top F). \tag{45}$$

Hence, with the adaptive weights fixed, the gradient of the original objective is

$$\begin{aligned} \nabla_F \mathcal{J}(F,\alpha,\beta) &= \sum_{v,g} \alpha_{v,g}^\gamma \nabla_F e_{v,g}(F) + \lambda \sum_g \beta_g^\gamma \nabla_F e_{*,g}(F) \\ &= 4 \sum_{v,g} \frac{\alpha_{v,g}^\gamma}{d_{v,g}} E_{v,g}(F)F + 4\lambda \sum_g \frac{\beta_g^\gamma}{d_{*,g}} \big[F(F^\top F) - Q_g(Q_g^\top F)\big]. \end{aligned} \tag{46}$$

After exact minimization over the adaptive weights, the reduced objective $\Phi(F)$ in Eq. 40 is smooth because all errors are positive. Using $p = 1/(\gamma - 1)$, $(1-\gamma)(-p) = 1$, and $p\gamma = p + 1$, direct differentiation of the reduced objective yields

$$\begin{aligned} \nabla\Phi(F) &= \sum_{v,g} \alpha_{v,g}(F)^\gamma \nabla_F e_{v,g}(F) + \lambda \sum_g \beta_g(F)^\gamma \nabla_F e_{*,g}(F) \\ &= \nabla_F \mathcal{J}(F, \alpha(F), \beta(F)). \end{aligned} \tag{47}$$

Thus, the gradient of the reduced objective is identical to the gradient of the original objective evaluated at the optimal adaptive weights.

A.5 Existence of a Global Minimizer

Proposition. Under the stated assumptions, the objective in Eq. (18) attains a global minimum on $\mathbb{R}^{N\times K} \times \mathcal{W}$.

Proof. Set $S = FF^\top$ and $d_{\max} = \max_g d_{*,g}$. Jensen's inequality gives $G^{1-\gamma} \le \sum_g \beta_g^\gamma \le 1$. Using $\|S - R\|_F^2 \ge \frac{1}{2}\|S\|_F^2 - \|R\|_F^2$ and $\|R^{(*,g)}\|_F^2 / d_{*,g} \le 1$, we obtain, uniformly over $\mathcal{W}$,

$$\mathcal{J}(F,\alpha,\beta) \ge \frac{\lambda G^{1-\gamma}}{2d_{\max}} \|FF^\top\|_F^2 - \lambda \ge \frac{\lambda G^{1-\gamma}}{2Kd_{\max}} \|F\|_F^4 - \lambda. \tag{48}$$

The second inequality follows from Cauchy–Schwarz applied to the eigenvalues of $F^\top F$, yielding $\|FF^\top\|_F^2 \ge \|F\|_F^4 / K$. Hence, $\mathcal{J}$ is coercive in F. Since $\mathcal{W}$ is compact and $\mathcal{J}$ is continuous, every nonempty sublevel set is compact. The extreme value theorem guarantees a global minimizer. □

B Experiment Settings

B.1 Datasets

Table 3 summarizes the statistics of the five benchmark datasets used in our experiments. Mfeat and MNIST-10k contain handwritten digit samples from ten categories. MSRC_v1 comprises images of different object categories, while Scene-15 includes images from fifteen indoor and outdoor scene categories. 100Leaves consists of leaf samples from one hundred plant species. Each dataset is represented by multiple feature views for clustering evaluation under incomplete view settings.

Table 3: Statistics of the benchmark datasets used in our experiments.

Dataset	Instances	Classes	Views	View dimensions
Mfeat	2,000	10	6	76/216/64/240/47/6
MSRC_v1	210	7	5	24/576/512/256/254
Scene-15	4,485	15	3	20/59/40
100Leaves	1,600	100	3	64/64/64
MNIST-10k	10,000	10	3	30/9/30

B.2 Incomplete View Generation

To evaluate AMGC under incomplete multi-view scenarios, missing views are simulated by randomly removing view observations from the original complete datasets. Following the common protocol in incomplete multi-view clustering, the missing rate r indicates the proportion of samples whose views are partially unavailable. Specifically, for a given missing rate r, $\lceil rN \rceil$ samples are randomly selected, and for each selected sample, one or more views are randomly removed while ensuring that at least one view remains available. The generated missing indicator matrix M is shared by all compared methods to ensure a fair comparison. In all experiments, the missing rates are set to $r \in \{0.1, 0.3, 0.5, 0.7\}$. In our evaluation protocol, missing views are randomly generated while the unselected samples retain complete observations. Therefore, the incomplete datasets generated preserve cross-view overlap in this specific setting.

B.3 Optimization Procedure

The proposed AMGC consists of three stages: relation construction, multi-granularity modeling, and consensus optimization. Anchors are selected based on sample completeness with FPS, and masked assignments are aggregated across views to construct the global relational reference, while local mutual k-nearest-neighbor graphs capture view-specific geometry. Local and global relations are then generated through multi-granularity modeling. The consensus embedding is initialized from the global relations and optimized with Adam by minimizing the weighted objective over local and global relations. The relation weights $\boldsymbol{\alpha}$ and global granularity weights $\boldsymbol{\beta}$ are updated according to their relation fitting errors. Final clustering is performed using k-means with 10 initializations on the consensus embedding F after training. Algorithm 1 summarizes the optimization procedure.

Algorithm 1 Optimization procedure of AMGC

Require: Incomplete multi-view data $\{X^{(v)}\}_{v=1}^{V}$, view masks M, anchor number m, granularities $\{c_g\}_{g=1}^{G}$, and maximum iterations T
Ensure: Consensus embedding F_{best}
1: Construct sample-anchor assignments based on the incomplete views.
2: Construct fixed local and global relations $\{R^{(v,g)}\}$ and $\{R^{(*,g)}\}$ via multi-granular modeling.
3: Initialize consensus embedding F from the global relations.
4: Initialize relation weights $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$, and set $J_{\text{best}} \leftarrow +\infty$.
5: **for** $t = 1$ to T **do**
6: Calculate relation fitting errors $e_{v,g}(F)$ and $e_{*,g}(F)$.
7: Update local relation weights $\boldsymbol{\alpha}$ and global granularity weights $\boldsymbol{\beta}$.
8: Update F using Adam and retain F_{best} at the lowest $\min_{\boldsymbol{\alpha},\boldsymbol{\beta}} J(F, \boldsymbol{\alpha}, \boldsymbol{\beta})$.
9: **end for**
10: Obtain final cluster labels by applying k-means to F_{best}.
11: **return** F_{best} and cluster labels.

B.4 Implementation Details

The proposed AMGC model is implemented in PyTorch and optimized using Adam. Each view is standardized using z-score normalization. The learning rate is set to 10^{-3} for 100Leaves and 10^{-2} for the other datasets. For Mfeat, MSRC_v1, Scene-15, 100Leaves, and MNIST-10k, the numbers

of anchors are 500, 48, 600, 600, and 800, respectively. In the same dataset order, the numbers of granularities are 8, 4, 8, 5, and 8, and the local neighborhood sizes are 8, 8, 5, 5, and 5. Granularity counts are generated by scaling K by G values evenly spaced from 0.5 to 2.0 and rounding. Each candidate and K itself are clamped to at most $\hat{m} = \min_{s \in \mathcal{S}} m_s$ (the fewest anchors available in any source); duplicates are removed, and the first G remaining values in ascending order yield $\{c_g\}_{g=1}^{G}$ (fewer when $\hat{m}$ is too small). The weight exponent γ and the global relation coefficient λ are set to 2 and 1, respectively. The training budgets are 1,000, 5,000, 6,000, 1,000, and 10,000 epochs in the same dataset order. The consensus embedding dimension is set to the number of clusters. Final clustering is performed using k-means with 10 initializations on the embedding at the lowest objective. All experiments are conducted on a workstation equipped with an Intel Core i9-13900K CPU (3.00 GHz), 128 GB RAM, and an NVIDIA GeForce RTX 4090 GPU with 24 GB memory.

C Supplementary Experiments

C.1 Loss Analysis

Figure 3 presents the evolution of the objective value and ACC on four datasets at a missing rate of 0.1. The objective values decrease rapidly during the early epochs and then gradually stabilize. Meanwhile, ACC generally improves despite initial fluctuations and reaches a relatively stable range in later epochs. Further reductions in the objective do not always produce corresponding accuracy gains, indicating that the clustering assignments can stabilize before the objective fully levels off. These results support the empirical stability of AMGC's optimization process.

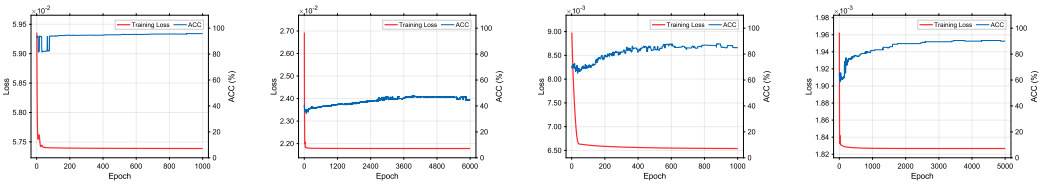

(a) Mfeat (b) Scene-15 (c) 100Leaves (d) MSRC_v1

Figure 3: Optimization trajectories of AMGC on four datasets at a missing rate of 0.1. The red and blue curves show the objective value (left axis) and ACC (%, right axis) over training epochs.

C.2 Visualization

Figure 4 compares the initial and optimized consensus embeddings on Mfeat and MNIST-10k at a missing rate of 0.7. After optimization, the embeddings exhibit clearer class separation and less mixing between different classes, particularly on MNIST-10k. Although some overlap remains, the improved class organization in the learned space provides qualitative evidence that AMGC learns more discriminative consensus representations even under substantial view missingness.

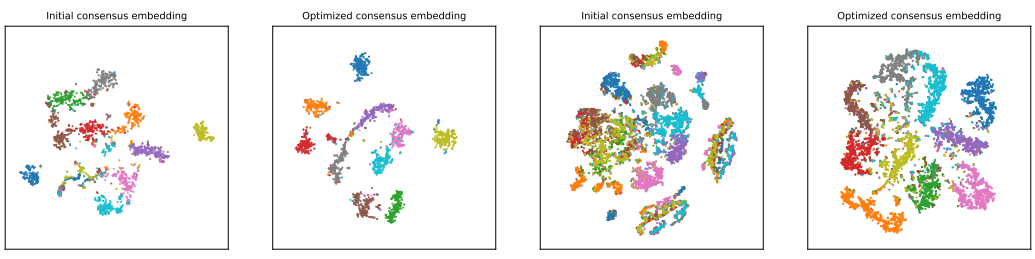

(a) Mfeat: Initial (b) Mfeat: Optimized (c) MNIST: Initial (d) MNIST: Optimized

Figure 4: t-SNE visualization of AMGC's initial and optimized consensus embeddings on Mfeat and MNIST-10k at a missing rate of 0.7. Points are colored by ground-truth class.

C.3 Complexity Analysis

Table 4 summarizes the computational complexity of the compared methods. AMGC has an overall complexity of $\mathcal{O}(N^2)$, mainly due to sample-level relation construction. Nevertheless, the anchor-based design performs structural modeling in a compact anchor space, and the matrix-free consensus optimization avoids explicitly storing dense relation matrices, reducing practical computational costs. At a missing-view rate of 0.3, Table 5 reports the actual computational costs. AMGC achieves lower optimization time and memory consumption on most datasets. These results demonstrate that AMGC maintains competitive efficiency while learning multi-granularity consensus structures through compact representations and efficient optimization with limited computational overhead.

Table 4: Complexity comparison between the comparative models and the proposed model.

Model	**Complexity**
DIMVC-HIA	$\mathcal{O}(N)$
DGIMVCM	$\mathcal{O}(N^3)$
DLPMAC	$\mathcal{O}(N^2)$
FreeCSL	$\mathcal{O}(N)$
RPCIC	$\mathcal{O}(N^2)$
AMGC (Ours)	$\mathcal{O}(N^2)$

Table 5: Computational efficiency comparison on five datasets at a missing-view rate of 0.3. Total time and epoch time are measured in seconds and milliseconds, respectively. Memory denotes peak tensor memory in MiB.

Datasets	Metrics	DIMVC-HIA	DGIMVCM	DLPMAC	FreeCSL	RPCIC	AMGC (Ours)
Mfeat	Total time (s)	1078.884	452.694	415.378	76.388	183.272	60.035
	Epoch time (ms)	10072.663	2262.584	146.989	389.455	326.505	30.205
	Memory (MiB)	792.321	2246.756	1337.658	636.696	372.369	95.666
MSRC_v1	Total time (s)	86.314	19.671	71.825	20.571	85.248	67.506
	Epoch time (ms)	821.223	195.425	16.764	136.386	37.991	13.041
	Memory (MiB)	398.041	1094.366	118.551	458.859	358.521	24.065
Scene-15	Total time (s)	1354.168	203.102	312.139	205.136	243.335	113.631
	Epoch time (ms)	12451.252	2029.862	297.501	1363.182	296.503	16.425
	Memory (MiB)	790.668	2288.805	1279.292	320.292	220.694	203.581
100Leaves	Total time (s)	2319.748	103.454	162.817	629.240	117.033	55.506
	Epoch time (ms)	20998.025	344.544	110.705	2846.306	105.312	25.213
	Memory (MiB)	900.952	997.714	822.671	326.815	180.897	81.436
MNIST-10k	Total time (s)	3274.098	1696.662	715.869	100.751	2938.904	215.588
	Epoch time (ms)	30896.488	8482.767	692.954	670.176	638.483	20.212
	Memory (MiB)	1408.416	6338.308	1329.605	319.118	295.705	820.735

C.4 Granularity-wise Weight Diagnostics

We examine the effective weight distribution across different granularities and the performance impact of adaptive global weighting. At a missing-view rate of 0.5, we compare AMGC with Uniform-β, which assigns equal β values across the same relation set, using five paired seeds. Table 6 shows that AMGC achieves higher mean ACC than Uniform-β on all five datasets, with improvements ranging from 0.45 to 6.70 percentage points. The improvement on Scene-15 is smaller than those on the other datasets, suggesting that the benefit of adaptive global weighting varies across datasets. Coarse relations receive a larger share of effective weight than their corresponding proportion in the relation set, while the remaining granularities still retain a substantial share. This distribution indicates that adaptive weighting adjusts the contributions of different granularities without concentrating exclusively on coarse relations. The effective weights reflect both the learned coefficients and normalization by squared relation norms, which jointly determine each granularity's contribution to the global fitting target.

Table 6: Adaptive weight analysis at a missing-view rate of 0.5. ACC is reported as mean $\pm$ standard deviation over five paired seeds. ΔACC is the difference between the mean ACC of AMGC and Uniform-β, measured in percentage points (pp). $G_<$ counts global granularities with $c_g < K$; $G_</G$ is their proportion in the relation set. $W_<$ is their total effective weight in the global fitting target, averaged over seeds.

	Clustering performance			Coarse-granularity mass	
Dataset	AMGC ACC (%)	Uniform-β ACC (%)	ΔACC (pp)	$G_</G$ (%)	$W_<$ (%)
Mfeat	91.35 ± 0.52	85.16 ± 1.07	$+6.19$	37.50	40.36
MNIST-10k	89.96 ± 0.04	83.52 ± 0.36	$+6.44$	37.50	39.87
Scene-15	41.75 ± 0.32	41.30 ± 2.23	$+0.45$	37.50	41.25
100Leaves	77.78 ± 0.45	71.08 ± 2.17	$+6.70$	40.00	50.92
MSRC-v1	76.19 ± 1.44	70.48 ± 3.16	$+5.71$	25.00	26.62

C.5 Analysis of Shared-Anchor Selection

To examine whether anchor selection introduces semantic bias, we conduct a controlled experiment on 100Leaves with a missing rate of 0.5 and an anchor budget of 100. We consider random missingness and class-dependent missingness, where view availability is correlated with class identity. All variants use the same missing masks, relation construction procedure, and optimization settings, differing only in the anchor selection strategy. Random selects anchors uniformly, Completeness selects samples with the most observed views, FPS applies farthest-point sampling, and AMGC combines completeness-based candidate filtering with FPS. In addition to clustering accuracy, we report three diagnostics of the selected anchor set $\mathcal{A}$, defined as

$$\mathrm{Cov}(\mathcal{A}) = \frac{|\{y_i : i \in \mathcal{A}\}|}{K},\ \mathrm{Avail}_{\min}(\mathcal{A}) = \min_v \frac{\sum_{i\in\mathcal{A}} M_{iv}}{|\mathcal{A}|},\ \rho_{\mathrm{miss}}(\mathcal{A}) = 1 - \frac{\sum_{i\in\mathcal{A}}\sum_{v=1}^{V} M_{iv}}{|\mathcal{A}|V}. \tag{49}$$

Here, $\mathrm{Cov}(\mathcal{A})$ measures the proportion of ground-truth classes represented by the selected anchors, $\mathrm{Avail}_{\min}(\mathcal{A})$ measures the minimum anchor availability across views, and $\rho_{\mathrm{miss}}(\mathcal{A})$ measures the proportion of unavailable anchor-view observations. Ground-truth labels are used only for post hoc evaluation of class coverage and are not involved in anchor selection.

Table 7: Analysis of shared-anchor selection under different missingness patterns on 100Leaves.

Missing condition	Metric	Random	Completeness	FPS	AMGC
random	ACC ↑	65.31 ± 1.20	60.99 ± 1.39	67.21 ± 0.93	71.58 ± 0.64
	Class coverage ↑	0.60 ± 0.0084	0.56 ± 0.0055	0.65 ± 0.0055	0.69 ± 0.0055
	Weakest-view availability ↑	0.7310 ± 0.0077	1.0000	0.9136 ± 0.0155	1.0000
	Unavailable ratio ↓	0.2266 ± 0.0045	0	0.0406 ± 0.0055	0
class_dep	ACC ↑	59.71 ± 1.32	53.30 ± 1.19	65.78 ± 0.94	69.78 ± 0.86
	Class coverage ↑	0.54 ± 0.0055	0.52 ± 0.0084	0.59 ± 0.0084	0.64 ± 0.0071
	Weakest-view availability ↑	0.7427 ± 0.0153	1.0000	0.9077 ± 0.0164	1.0000
	Unavailable ratio ↓	0.2460 ± 0.0113	0	0.0742 ± 0.0064	0

As shown in Table 7, completeness-based selection improves cross-view availability but yields lower class coverage, as it prioritizes samples observed in more views. In contrast, FPS promotes feature-space diversity, while combining completeness screening with FPS achieves broader class coverage while maintaining full cross-view availability of the selected anchors under both missingness patterns. Under class-dependent missingness, AMGC consistently achieves higher clustering accuracy and class coverage than the alternative selection strategies, demonstrating the effectiveness of jointly considering observation availability and feature-space diversity. These results suggest that shared-anchor construction benefits from balancing observation completeness and feature-space diversity in incomplete multi-view settings.