FROM MISSING RELATIONS TO RELIABLE STRUCTURE: CROSS-VIEW RELATION RECOVERY AND LATENT STRUCTURE ESTIMATION FOR INCOMPLETE MULTI-VIEW LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

In real-world applications, multi-view data are often incomplete because some views are unavailable for some samples. Although existing incomplete multi-view learning methods have made substantial progress in feature recovery and structural modeling, view missingness can also cause relational information to disappear when relation endpoints are unavailable, leaving the underlying structure incomplete. To address this issue, we propose **Reliable Latent Structure Estimation (ReLaS)**, a relation-centric framework that explicitly restores missing relational structure and learns a latent consensus graph from incomplete multi-view observations. For relations that remain directly observable, ReLaS estimates their relation evidence from endpoint availability and cross-view edge confirmations, while characterizing their reliability. Relations rendered unobservable by missing endpoints are explicitly recovered through cross-view structural correspondence, with the reliability of recovered relations further calibrated in a self-supervised manner. Both direct and recovered relation evidence are then integrated under relation-profile consistency to obtain the latent consensus graph. Experiments on multiple benchmark datasets under different view-missing rates demonstrate the effectiveness of ReLaS in restoring relational structure and improving semi-supervised classification performance.

1 INTRODUCTION

Multi-view data arise naturally when the same objects are described by multiple sources, modalities, or feature sets, providing complementary and consistent information from different perspectives (Xu et al., 2013; Chao et al., 2021). Multi-view learning aims to integrate such heterogeneous information while preserving cross-view consensus and complementarity, leading to a variety of paradigms including kernel-based learning, subspace learning, graph-based modeling, and deep representation learning (Chao et al., 2021; Chen et al., 2022; Haris et al., 2024). Accordingly, multi-view learning has been widely studied across clustering, classification, and semi-supervised learning, and remains an important framework for learning from complex multi-source data (Yu et al., 2025).

In practice, complete observations across all views are often unavailable because of equipment failure, data transmission issues, or other acquisition limitations, giving rise to incomplete multi-view learning (IMVL) (Tang et al., 2024). Existing approaches can be broadly divided into two major directions, feature recovery and structure learning. Feature-recovery methods seek to reconstruct missing views or learn more complete representations from the available observations. Liu et al. (2024a) use cross-view reconstruction to alleviate anchor shift, while ARSL-IMVC jointly performs collaborative view recovery and tensorial subspace learning through latent representations (Wang et al., 2026b). Diffusion Contrastive Generation combines diffusion-based recovery with contrastive learning (Zhang et al., 2025), whereas Wu et al. (2025a) proposed I2MVC to avoid explicit view imputation through pseudo-supervised knowledge distillation. The other major direction focuses on structure learning, where sample similarities, graph structures, or pairwise relations are directly modeled from incomplete observations (Bai et al., 2024). Liu et al. (2024b) proposed SCSL to learn complete cross-view similarities from intra-view and inter-view sample relations, and HCLS-CGL

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.

exploits highly confident local structures to guide consensus graph learning (Wen et al., 2023). AGF-TI further recovers incomplete graph structures through adversarial graph fusion and tensor learning (Jiang et al., 2025b), while RCIMVC takes a relation-level perspective by directly completing relation graph features rather than reconstructing the original missing data (Wu et al., 2025b). However, existing structural methods largely focus on constructing, refining, or recovering relational structures from incomplete observations, while the relational incompleteness induced directly by view missingness has not been systematically characterized and restored at the relation level.

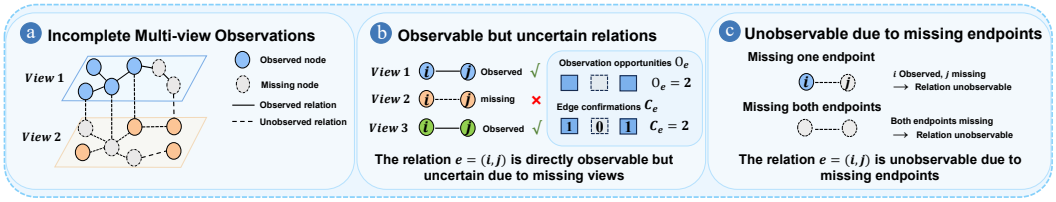

Figure 1: Illustration of relational incompleteness caused by view missingness.

As illustrated in Fig. 1, view missingness leads to two relation states. When both endpoints are observed, the relation can still be directly assessed, although the amount of supporting evidence may vary across views. Furthermore, when an endpoint is missing, the relation becomes unobservable in that view and disappears from the corresponding view-specific graph, resulting in relational incompleteness. This suggests that incomplete multi-view data should be considered not only in terms of feature incompleteness, but also at the relation level, where missing views directly alter the availability of pairwise structural evidence. Existing structural methods mainly focus on learning from or refining observable relations under incomplete observations, while some further explore structural recovery for specific missing patterns. Nevertheless, relation loss caused specifically by missing endpoints is rarely isolated and explicitly treated as part of a unified structural restoration problem. Therefore, a more complete treatment of relational structure should jointly consider both relation states by reliably estimating directly observable relations, explicitly recovering relations rendered unobservable by missing endpoints, and integrating both into a latent consensus structure.

To address this challenge, we propose Reliable Latent Structure Estimation (ReLaS), a relation-centric framework that explicitly restores missing relational structure and learns a reliable latent consensus graph from incomplete multi-view observations. The overall framework of ReLaS is illustrated in Fig. 2. ReLaS first performs reliability-aware estimation of directly observable relations by combining endpoint-derived observation opportunities with cross-view edge confirmations. It then recovers relations rendered unobservable by missing endpoints through cross-view structural correspondence and further calibrates the reliability of the recovered relations in a self-supervised manner. Finally, reliable direct and recovered relation evidence are integrated under relation-profile consistency to obtain the latent consensus graph. The main contributions of this study are summarized as follows:

- A reliability-aware direct relation estimation mechanism is developed to estimate directly observable relations from endpoint availability and cross-view confirmations, with relation-specific reliability explicitly characterized.
- A cross-view relation recovery mechanism is introduced to explicitly recover relations rendered unobservable by missing endpoints through cross-view structural correspondence, with the reliability of recovered relations further calibrated in a self-supervised manner.
- A latent consensus graph learning formulation is developed to integrate reliable direct and recovered relation evidence under relation-profile consistency, thereby restoring relational structure from incomplete multi-view observations.

2 Related Work

Existing incomplete multi-view learning methods can be broadly organized into two major directions, feature recovery and structural learning. Feature recovery methods seek to reconstruct missing observations or infer more complete latent representations from the available views. Chao et al. (2024) proposed ICMVC to combine relation transfer with graph convolution and contrastive representation

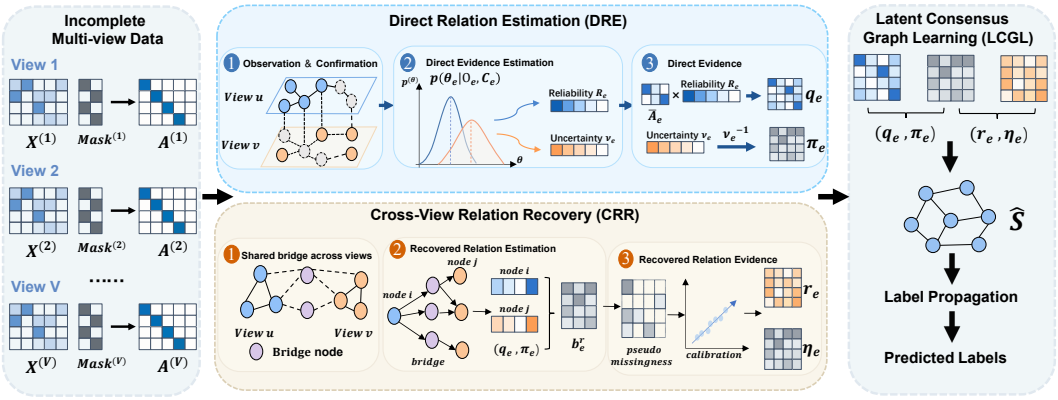

Figure 2: Illustration of the proposed ReLaS framework. ReLaS addresses relational incompleteness through Direct Relation Estimation (DRE) and Cross-View Relation Recovery (CRR), which estimate reliable direct relations (q_e, π_e) and recover missing relations with calibrated reliability (r_e, η_e), respectively. Latent Consensus Graph Learning (LCGL) then integrates both sources of relation evidence to obtain the latent consensus graph $\widehat{S}$ for subsequent label propagation.

learning. UNIFIER jointly performs view imputation and feature selection under sample- and feature-level similarity guidance, while PCL-UIMVC and UGCF further couple missing-information recovery with cross-view reconstruction or graph-based learning (Huang et al., 2024; Cai et al., 2024; Wang et al., 2024). More recent studies place greater emphasis on the quality of recovered representations. LRGR refines graph-imputed representations to alleviate distribution shift and representation homogenization, and HSACC combines hierarchical semantic alignment with cooperative completion (Xi et al., 2025; Ding et al., 2025). Other approaches explore hierarchical recovery, imputation-free latent learning, or cross-view generative modeling (Du et al., 2026; Chen et al., 2026; Gao et al., 2026). Overall, feature-level methods have moved beyond direct imputation toward richer representation recovery, but their primary focus remains on the representation space rather than relational structure itself.

Structural learning methods directly model sample similarities, graph structures, and relational information under view missingness. Early studies mainly learn or refine structural representations from the available observations. Xing et al. (2024) proposed RTRGC to recover incomplete similarity graphs through tensor-ring completion and robust estimation, while ASCR jointly learns and completes anchor-sample similarities together with latent representations (Xu et al., 2025). SFCR learns a unified similarity graph to guide missing-sample recovery, and CAL exploits sample-wise cross-view affinities and high-order interactions to construct a unified structural graph (Jiang et al., 2025a; Liu et al., 2025). The reliability of learned structures has also received increasing attention, as reflected by HLSCG-IMC and SWCGLIMVC (Zhao et al., 2025; Sun et al., 2025). As research has progressed, more attention has been paid to structural recovery under view missingness. However, even among structure-oriented methods, recovery is often coupled with missing-sample or latent-representation reconstruction, followed by structure learning or refinement (Xu et al., 2025; Jiang et al., 2025a; Wang et al., 2024). More recent methods, including AGF-TI, RCIMVC, FGIMC, and AGLF, begin to restore missing structural information more directly (Jiang et al., 2025b; Wu et al., 2025b; Ding et al., 2026; Wang et al., 2026a). Taken together, although structural recovery has begun to receive increasing attention, many existing structural methods still build upon feature recovery, whereas direct relation-level restoration remains comparatively less systematically studied in incomplete multi-view learning.

3 METHOD

Let $\{X^{(v)}\}_{v=1}^{V}$ denote n aligned samples from V views, where $X^{(v)} \in \mathbb{R}^{n \times d_v}$, and let $M \in \{0,1\}^{n \times V}$ be the view-availability mask. For each view, a weighted sparse affinity graph $A^{(v)}$ is constructed using only its observed samples. ReLaS models incomplete multi-view data at the

relation level by combining reliable direct relation estimation, cross-view relation recovery, and latent consensus graph learning.

3.1 Direct Relation Estimation (DRE)

For an unordered relation $e = (i, j)$, define $O_e^{(v)} = M_{iv}M_{jv}$ and $C_e^{(v)} = O_e^{(v)}\mathbb{I}[A_{ij}^{(v)} > 0]$ as its observation-opportunity and edge-confirmation indicators in view v, respectively. Aggregating them across views gives

$$O_e = \sum_{v=1}^{V} O_e^{(v)}, \qquad C_e = \sum_{v=1}^{V} C_e^{(v)}. \tag{1}$$

Here, O_e counts the views in which both endpoints are jointly observed and hence the edge can be assessed, whereas C_e counts the views that actually confirm the edge, separating observation opportunities from positive edge evidence.

Let $\mathcal{E}_D = \{e : C_e > 0\}$ denote the direct relation set, and let $\bar{A}_e = \frac{1}{C_e}\sum_{v:C_e^{(v)}=1} A_{ij}^{(v)}$ denote the mean edge weight over the views that confirm e. To account for both the frequency of edge confirmations and the amount of available evidence, ReLaS models the latent opportunity-conditioned edge-confirmation probability θ_e with the non-informative prior $\theta_e \sim \mathrm{Beta}(1, 1)$. Conditional on θ_e, each observable view provides a Bernoulli confirmation

$$C_e^{(v)} \mid O_e^{(v)} = 1, \theta_e \sim \mathrm{Bernoulli}(\theta_e). \tag{2}$$

Combining the prior with the observed confirmations yields

$$\theta_e \mid C_e, O_e \sim \mathrm{Beta}\left(C_e + 1,\, O_e - C_e + 1\right). \tag{3}$$

This Beta–Bernoulli construction provides a working posterior model for aggregating cross-view confirmation evidence. DRE then defines the conservative confirmation reliability using its lower δ-quantile

$$R_e = Q_\delta[\mathrm{Beta}\left(C_e + 1,\, O_e - C_e + 1\right)], \qquad \delta = 0.05. \tag{4}$$

Combining this reliability with the observed relation magnitude gives the direct relation center

$$q_e = \bar{A}_e R_e. \tag{5}$$

Thus, q_e jointly reflects the observed affinity magnitude and the reliability of its cross-view confirmations.

The same posterior further characterizes relation-specific uncertainty:

$$\nu_e = \mathrm{Var}(\theta_e \mid C_e, O_e) = \frac{(C_e + 1)(O_e - C_e + 1)}{(O_e + 2)^2(O_e + 3)}. \tag{6}$$

The corresponding relation precision is defined as

$$\pi_e = \nu_e^{-1}. \tag{7}$$

A larger π_e indicates higher confidence in the corresponding relation center. Before optimization, the precision values are robustly rescaled to limit the influence of extreme values.

3.2 Cross-View Relation Recovery (CRR)

Relations involving a missing endpoint cannot be directly assessed in that view; CRR recovers them through structural correspondence across overlapping views. Let $\Omega_v = \{i : M_{iv} = 1\}$ denote the observed nodes in view v, and $\mathcal{B}_{uv} = \Omega_u \cap \Omega_v$ the bridge shared by views u and v. For an ordered view pair $u \to v$, consider a node $i \in \Omega_u \setminus \Omega_v$, which is observed in the source view but missing from the target view.

CRR first identifies plausible relations for i by transferring local structure through the shared bridge. Specifically, let $\mathcal{A}_{u\to v}(i)$ contain the K_A bridge nodes most strongly connected to i in the source view u. Their neighborhoods in the target view v are then transferred to i:

$$\mathcal{C}_{u\to v}(i) = \bigcup_{a \in \mathcal{A}_{u\to v}(i)} \mathcal{N}_v(a), \tag{8}$$

where $\mathcal{N}_v(a)$ denotes the nodes connected to a in the target-view affinity graph $A^{(v)}$. Candidate pairs are collected over all ordered view pairs, deduplicated, and those already contained in $\mathcal{E}_D$ are excluded, yielding $\mathcal{E}_{\mathrm{cand}}$.

To characterize a node relative to the shared bridge, CRR uses its affinities to all bridge nodes. The normalized bridge-relative relation profile of i is

$$\boldsymbol{\phi}_i^{u|uv} = \frac{\left[e^{-d_u(i,a)/\sigma_u}\right]_{a\in\mathcal{B}_{uv}}}{\left\|\left[e^{-d_u(i,a)/\sigma_u}\right]_{a\in\mathcal{B}_{uv}}\right\|_2}. \tag{9}$$

For the ordered view pair $u\to v$, let $\gamma=(u\to v)$ denote the corresponding cross-view channel. For candidate $e=(i,j)$, its structural correspondence in this channel is measured by

$$b_e^{\gamma} = \left\langle \boldsymbol{\phi}_i^{u|uv}, \boldsymbol{\phi}_j^{v|uv} \right\rangle. \tag{10}$$

Because b_e^{γ} measures bridge-relative structural similarity rather than relation strength itself, CRR maps it onto the direct-relation scale learned by DRE. For the direct relations $\mathcal{D}_\gamma$ available in channel γ, the mapping is estimated by

$$(\alpha_\gamma,\beta_\gamma) = \underset{\alpha\in\mathbb{R},\,\beta\ge 0}{\arg\min} \sum_{f\in\mathcal{D}_\gamma} \pi_f \left(q_f - \alpha - \beta b_f^{\gamma}\right)^2. \tag{11}$$

This yields the calibrated channel response $g_\gamma(b)=\mathrm{clip}(\alpha_\gamma+\beta_\gamma b,0,1)$. Finally, aggregating the valid channels Γ_e yields the recovered relation center

$$r_e = \frac{1}{|\Gamma_e|}\sum_{\gamma\in\Gamma_e} g_\gamma(b_e^{\gamma}). \tag{12}$$

The recovered center r_e quantifies relation strength but not its reliability. CRR therefore constructs a self-supervised reliability signal through controlled pseudo missingness.

Specifically, one observed view is deterministically hidden from each eligible multi-view sample, yielding a pseudo-incomplete observation pattern. To ensure that the pseudo target reflects relation loss caused by the imposed missingness rather than changes from graph reconstruction, the edge confirmations $C_e^{(v)}$ induced by the original affinity graphs $A^{(v)}$ are kept fixed, while DRE and CRR are rerun using only the remaining observations. This produces a pseudo-recovered center r_e^{ps} for each available candidate. Its calibration target is defined by

$$y_e^{\mathrm{ps}} = \mathbb{I}\left[e\in\mathcal{E}_D \ \wedge \sum_{v=1}^{V}\widetilde{O}_e^{(v)}C_e^{(v)}=0\right], \tag{13}$$

where $\widetilde{O}_e^{(v)}$ denotes the observation opportunity after pseudo masking. Hence, $y_e^{\mathrm{ps}}=1$ identifies a relation that was directly confirmed before masking but loses all such confirmations under the imposed missingness.

Collecting the resulting pairs gives $\mathcal{D}_{\mathrm{ps}}=\{(r_\ell^{\mathrm{ps}},y_\ell^{\mathrm{ps}})\}_{\ell=1}^{N_{\mathrm{ps}}}$. Isotonic regression is then used to calibrate recovered relation strength against the pseudo-recovery outcomes:

$$h^\star = \underset{h\in\mathcal{H}_\uparrow}{\arg\min}\sum_{\ell=1}^{N_{\mathrm{ps}}}\left(y_\ell^{\mathrm{ps}} - h(r_\ell^{\mathrm{ps}})\right)^2, \tag{14}$$

where $\mathcal{H}_\uparrow$ contains nondecreasing mappings to $[0,1]$.

Let $\bar{y}=N_{\mathrm{ps}}^{-1}\sum_\ell y_\ell^{\mathrm{ps}}$ denote the pseudo-recovery base rate, i.e., the empirical fraction of positive recovery targets in $\mathcal{D}_{\mathrm{ps}}$. The calibrated reliability associated with the recovered relation center r_e is then quantified by

$$\eta_e = \left[\frac{h^\star(r_e)-\bar{y}}{1-\bar{y}}\right]_+. \tag{15}$$

Thus, η_e measures the calibrated reliability of r_e beyond the pseudo-recovery base rate. Candidates with $\eta_e>0$ form $\mathcal{E}_R=\{e\in\mathcal{E}_{\mathrm{cand}}:\eta_e>0\}$.

3.3 Latent Consensus Graph Learning (LCGL)

DRE provides (q_e, π_e) for direct relations in $\mathcal{E}_D$, while CRR provides (r_e, η_e) for recovered relations in $\mathcal{E}_R$. LCGL integrates both sources over the expanded relation set $\mathcal{E}^+ = \mathcal{E}_D \cup \mathcal{E}_R$.

Let Q be the sparse symmetric graph supported on $\mathcal{E}_D$, with $Q_{ij} = q_e$ for $e = (i, j)$, and let $\mathcal{L}_Q = I - D_Q^{-1/2} Q D_Q^{-1/2}$ denote its normalized Laplacian. Combining direct- and recovered-relation fidelity with structural regularization yields the latent consensus graph

$$\widehat{S} = \arg\min_{S \in \mathcal{G}_{\mathcal{E}^+}} \frac{1}{2} \sum_{e \in \mathcal{E}_D} \pi_e (S_e - q_e)^2 + \frac{1}{2} \sum_{e \in \mathcal{E}_R} \eta_e (S_e - r_e)^2 + \frac{\lambda}{2} \operatorname{tr}(S^\top \mathcal{L}_Q S), \tag{16}$$

where $S_e = S_{ij}$ for $e = (i, j)$ and $\mathcal{G}_{\mathcal{E}^+} = \{S : S = S^\top,\ S \geq 0,\ \operatorname{diag}(S) = 0,\ S_{ij} = 0 \text{ for } (i, j) \notin \mathcal{E}^+\}$. The first term anchors direct relations to q_e according to their precision π_e, while the second anchors recovered relations to r_e according to their calibrated reliability η_e. The third term regularizes the global relation structure.

Since each row $S_{i,:}$ represents the relation profile of sample i, the structural term can be written as

$$\operatorname{tr}(S^\top \mathcal{L}_Q S) = \frac{1}{2} \sum_{i,j} Q_{ij} \left\| \frac{S_{i,:}}{\sqrt{d_i^Q}} - \frac{S_{j,:}}{\sqrt{d_j^Q}} \right\|_2^2, \tag{17}$$

where $d_i^Q = (D_Q)_{ii}$. Hence, samples strongly connected in Q are encouraged to exhibit compatible relation profiles in the latent consensus graph. Since Q contains only direct relations, direct evidence determines the regularization geometry, while $\mathcal{E}_R$ enlarges the admissible relation set.

The structural coefficient is fixed to $\lambda = 0.5$ based on curvature-scale matching between the direct-relation fidelity and structural regularization terms, as detailed in Appendix A.5.

For numerical optimization, each undirected relation in $\mathcal{E}^+$ is represented by a single nonnegative edge variable, which enforces symmetry and avoids optimizing a dense $n \times n$ matrix. The resulting bound-constrained quadratic problem is solved using L-BFGS-B. Further optimization details and the convexity analysis are provided in the appendix.

For semi-supervised prediction, labels are propagated on the learned latent consensus graph through the standard regularized graph readout $F = (I + \tau \mathcal{L}_{\widehat{S}})^{-1} Y$. Labels are introduced only after latent consensus graph estimation.

4 Experiments

4.1 Experimental Setting

Datasets: Experiments are conducted on eight benchmark datasets, including CUB (Li et al., 2023), Caltech20 (Lin et al., 2023), COIL20 (Hu et al., 2025), Scene15 (Lin et al., 2023), ALOI-100 (Dai et al., 2025), Reuters (Jin et al., 2025), LandUse-21 (Lin et al., 2023), and Fashion (Zhang et al., 2025). **Competitors:** ReLaS is compared with eight representative methods, including SLIM (Yang et al., 2018), LHGN (Zhu et al., 2023), DCP (Lin et al., 2023), AMSC (Zhuge et al., 2024), AGF-TI (Jiang et al., 2025b), MEGNN (Wang et al., 2025), CoGFormer (Lai et al., 2026), and MHML (Wang et al., 2027). For methods requiring complete data, we follow AGF-TI and recover missing features using Deep Matrix Factorization (DMF) before evaluation. **Incomplete Data Construction:** Different levels of view missingness are considered by varying the VMR from 0.1 to 0.8, while 10% of the samples are provided with labels. **Evaluation Metrics:** Classification performance is evaluated using Accuracy (ACC) and Macro-F1.

4.2 Semi-Supervised Classification

Table 1 shows that ReLaS achieves strong classification performance under different levels of view missingness. While several competing methods remain competitive on individual datasets at relatively low VMRs, the advantage of ReLaS becomes more evident as the missingness increases. At VMRs

Table 1: Classification performance (%) under different view-missing rates (VMRs) with 10% labeled samples. Results for each method are reported as ACC / Macro-F1. The best result is shown in **bold**, and the second-best result is underlined.

VMR	Method	CUB	Caltech20	COIL20	Scene15	ALOI-100	Reuters	LandUse-21	Fashion
0.3	SLIM	64.72 / 63.95	42.41 / 6.02	52.11 / 47.88	35.28 / 33.02	OOM	53.89 / 53.70	20.03 / 16.26	OOM
	LHGN	71.67 / 71.42	83.69 / 65.22	93.25 / 93.20	62.96 / 61.60	61.50 / 61.49	55.60 / 40.63	17.80 / 17.16	87.54 / 83.60
	DCP	65.19 / 64.57	76.99 / 54.91	85.04 / 84.64	66.96 / 64.91	69.90 / 69.60	27.03 / 24.24	37.43 / 35.36	85.24 / 83.66
	AMSC	72.41 / 71.22	52.10 / 10.75	84.45 / 82.88	59.28 / 54.26	69.54 / 68.21	33.70 / 31.53	33.54 / 30.23	66.71 / 61.76
	AGF-TI	78.33 / 78.30	81.66 / 60.16	87.34 / 87.22	66.85 / 63.99	83.23 / 82.38	55.97 / 56.16	34.79 / 32.28	87.82 / 87.31
	MEGNN	61.96 / 60.73	70.05 / 43.47	76.76 / 74.96	48.62 / 44.63	46.13 / 44.93	39.26 / 34.24	37.33 / 36.26	85.52 / 85.24
	CoGFormer	75.19 / 74.96	83.13 / 65.33	92.38 / 92.37	65.13 / 63.95	81.35 / 81.87	57.50 / 57.37	42.25 / **41.80**	**90.61 / 90.57**
	MHML	73.15 / 72.63	81.92 / 60.39	90.23 / 90.12	58.98 / 56.74	50.75 / 52.58	57.31 / 57.19	33.57 / 32.80	89.42 / 89.32
	ReLaS	**80.22 / 79.75**	**84.04 / 66.78**	**93.55 / 93.53**	**70.78 / 69.37**	**90.63 / 90.62**	**60.19 / 59.73**	**42.69** / 41.41	90.17 / 90.01
0.5	SLIM	56.78 / 55.94	43.39 / 6.93	46.09 / 42.32	32.47 / 30.08	OOM	52.43 / 52.28	18.29 / 14.49	OOM
	LHGN	63.20 / 61.51	81.97 / 62.83	91.32 / 91.24	60.15 / 59.07	51.82 / 52.02	55.48 / 42.23	15.86 / 14.26	84.68 / 84.00
	DCP	58.02 / 56.50	70.54 / 47.60	80.74 / 79.86	59.09 / 57.30	59.07 / 59.45	23.36 / 16.33	34.28 / 33.38	81.31 / 80.95
	AMSC	65.63 / 63.36	54.64 / 13.23	85.08 / 83.67	57.57 / 52.95	63.86 / 62.61	28.87 / 26.10	32.02 / 29.69	63.30 / 58.30
	AGF-TI	73.28 / 71.40	80.80 / 57.06	87.18 / 86.97	64.90 / 61.09	82.31 / 81.59	54.75 / 54.69	36.60 / 34.59	86.00 / 85.32
	MEGNN	58.33 / 56.31	62.37 / 41.48	69.94 / 68.68	41.57 / 39.88	37.52 / 34.23	37.30 / 34.00	30.64 / 29.26	83.13 / 83.38
	CoGFormer	73.48 / 73.12	81.60 / 61.44	89.31 / 89.20	59.52 / 58.08	76.19 / 77.31	54.91 / 55.06	37.35 / 37.00	87.79 / 87.65
	MHML	69.48 / 69.12	80.15 / 58.09	88.11 / 87.89	53.58 / 51.97	42.71 / 45.73	56.92 / 56.96	29.28 / 28.54	87.64 / 87.58
	ReLaS	**77.13 / 76.56**	**82.57 / 64.04**	**92.33 / 92.33**	**69.50 / 67.99**	**89.34 / 89.33**	**60.03 / 59.53**	**39.98 / 38.75**	**88.26 / 88.01**
0.7	SLIM	53.15 / 52.91	43.81 / 7.64	36.48 / 34.01	27.02 / 24.55	OOM	50.00 / 49.63	13.41 / 10.93	OOM
	LHGN	58.15 / 57.73	79.23 / 60.10	90.24 / 90.12	57.08 / 55.90	41.21 / 41.45	52.82 / 35.23	12.83 / 11.66	82.06 / 82.13
	DCP	55.19 / 53.83	65.91 / 42.95	70.23 / 67.91	52.53 / 50.15	44.08 / 41.92	22.87 / 15.74	28.70 / 27.05	75.03 / 73.03
	AMSC	62.41 / 61.46	56.14 / 15.45	81.33 / 79.45	52.99 / 49.95	56.05 / 54.43	25.56 / 20.69	28.62 / 26.62	59.28 / 54.14
	AGF-TI	70.37 / 63.70	79.49 / 55.17	85.62 / 85.10	63.71 / 59.49	79.02 / 78.97	52.87 / 52.81	35.71 / 34.38	84.36 / 83.68
	MEGNN	55.93 / 53.05	51.66 / 34.37	63.67 / 60.64	37.45 / 32.98	31.52 / 31.45	36.20 / 30.37	27.22 / 26.06	80.22 / 79.56
	CoGFormer	68.06 / 67.40	80.47 / 58.95	83.67 / 83.63	54.53 / 53.43	69.47 / 70.94	51.85 / 51.29	33.36 / 32.91	86.16 / 86.04
	MHML	67.59 / 67.21	79.07 / 56.33	83.71 / 83.50	49.29 / 46.92	36.95 / 38.81	52.22 / 52.05	26.27 / 25.48	86.08 / 85.85
	ReLaS	**75.53 / 75.11**	**81.46 / 62.03**	**90.82 / 90.74**	**67.35 / 65.63**	**87.88 / 87.86**	**59.44 / 58.73**	**37.77 / 36.42**	**86.38 / 86.12**

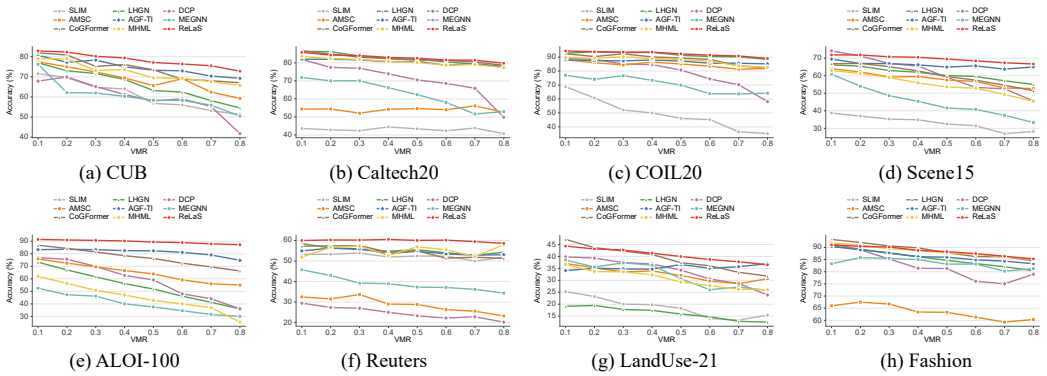

(a) CUB (b) Caltech20 (c) COIL20 (d) Scene15

(e) ALOI-100 (f) Reuters (g) LandUse-21 (h) Fashion

Figure 3: Classification accuracy under different view-missing rates on all eight benchmark datasets.

of 0.5 and 0.7, ReLaS consistently shows strong competitiveness across all datasets, indicating that its relation modeling strategy remains effective under increasingly degraded observations.

Fig. 3 further reports the classification accuracy over VMRs from 0.1 to 0.8. ReLaS maintains stable performance across the entire range, and on datasets such as Caltech20, LandUse-21, and Fashion, its relative advantage becomes more pronounced at higher VMRs even when it is not always the best under mild missingness. This trend is consistent with the complementary roles of DRE and CRR, where DRE preserves reliable evidence from relations that remain directly observable, while CRR explicitly restores relational information that becomes progressively unavailable as view missingness increases, thereby helping ReLaS preserve more informative relational structure under severe degradation.

4.3 Ablation Study

To evaluate the contribution of each component, we conduct an ablation study by progressively introducing LCGL, DRE, and CRR into the raw observed relation graph, as indicated by the checkmarks

Table 2: Ablation results (%) with VMR = 0.5 and 10% labeled samples. Results are reported as ACC / Macro-F1.

LCGL	DRE	CRR	CUB	Caltech20	COIL20	Scene15	ALOI-100	Reuters	LandUse-21	Fashion
×	×	×	70.37/69.92	79.44/57.78	89.29/89.25	66.31/65.57	86.60/86.56	55.64/55.37	35.87/34.99	86.58/86.33
✓	×	×	71.91/71.44	80.14/58.85	89.61/89.57	67.25/65.84	87.71/87.68	55.74/55.24	36.40/35.43	86.82/86.67
✓	✓	×	74.07/73.73	81.82/61.85	91.48/91.27	68.57/67.68	88.97/88.96	57.40/57.12	37.09/36.24	86.92/86.76
✓	✓	✓	**77.13/76.56**	**82.57/64.04**	**92.33/92.33**	**69.50/67.99**	**89.34/89.33**	**60.03/59.53**	**39.98/38.75**	**88.26/88.01**

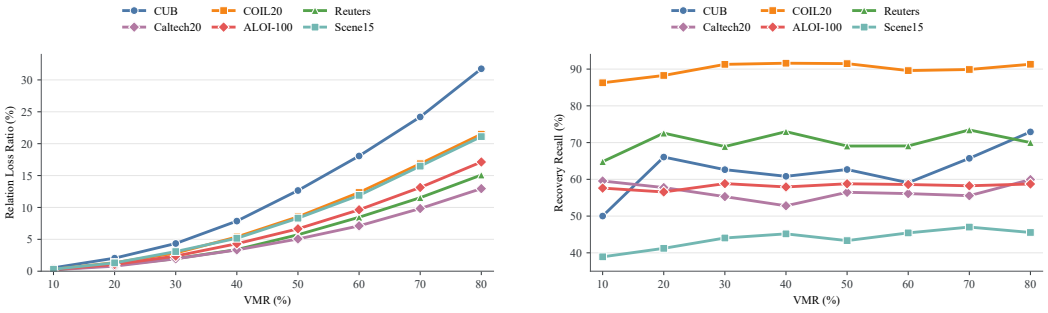

(a) Relation loss ratio under different VMRs.

(b) Recovery recall of CRR under different VMRs.

Figure 4: Relation loss and recovery analysis under different view-missing rates.

in Table 2. Introducing LCGL alone generally brings only limited improvements, suggesting that structural refinement without additional reliable relation evidence has a relatively modest effect. With DRE further incorporated, the performance gains become more evident on most datasets, showing the benefit of reliable direct relation estimation under view missingness. After CRR is introduced to recover relations that become unobservable because of missing views, the complete ReLaS model achieves the best overall performance, demonstrating the complementary contributions of reliable direct relation estimation, missing relation recovery, and latent consensus graph learning.

4.4 Relation Degradation and Recovery Analysis

To investigate the relation degradation caused by view missingness and the recovery capability of CRR, we analyze the relation loss ratio and recovery recall under different VMRs. The relation loss ratio is computed as the proportion of reference relations that become completely unobservable because no view jointly observes both endpoints. Recovery recall is computed as the proportion of these lost relations that are successfully recovered and retained by CRR after reliability calibration. We select several representative datasets for visualization, and the results are shown in Fig. 4.

As shown in Fig. 4a, the relation loss ratio increases consistently as the VMR grows across different datasets, indicating that view missingness leads to increasingly severe relational information loss. This observation demonstrates that incomplete views can affect not only feature availability but also the preservation of relational structures. Fig. 4b further illustrates the recovery recall of CRR on the lost relations. CRR is able to recover a considerable proportion of missing relations under different missing rates, while the recovery performance varies across datasets due to differences in cross-view structural consistency. For example, COIL20 maintains a relatively high recovery recall across various VMRs, whereas other datasets exhibit different recovery patterns. These results verify that CRR can effectively exploit complementary structural information across views for missing relation recovery.

4.5 Parameter Sensitivity

To evaluate the parameter sensitivity of ReLaS, the key parameter λ is selected for analysis. Fig. 5 reports the Accuracy and Macro-F1 results on all eight benchmark datasets under different values of λ. ReLaS maintains stable performance across a broad range of λ, demonstrating low sensitivity to

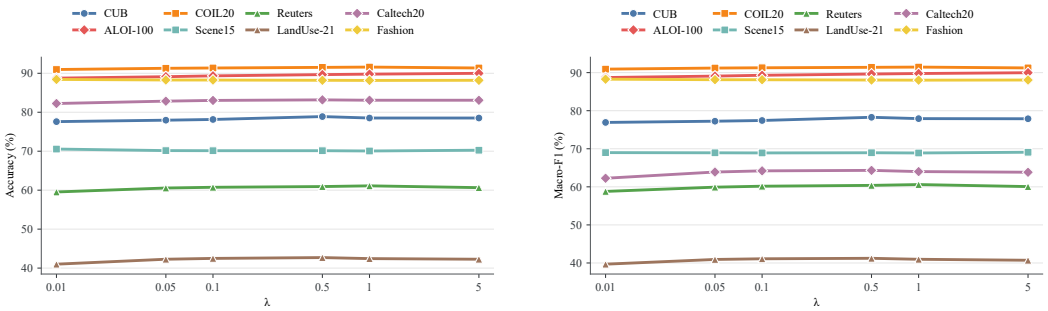

(a) Accuracy under different values of λ.

(b) Macro-F1 under different values of λ.

Figure 5: Parameter sensitivity of ReLaS with respect to the relation-profile consistency weight λ.

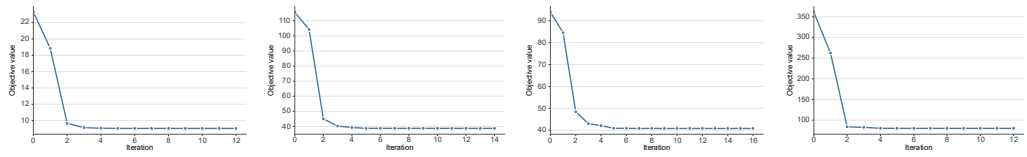

(a) CUB (b) Caltech20 (c) COIL20 (d) Scene15

(e) ALOI-100 (f) Reuters (g) LandUse-21 (h) Fashion

Figure 6: Convergence behavior of the LCGL optimization on all benchmark datasets.

this parameter. The consistent performance across different datasets further suggests that λ does not require careful tuning in practice.

4.6 Convergence Analysis

To examine the convergence of ReLaS, the objective value at each iteration of the LCGL optimization is recorded. As shown in Fig. 6, the objective value decreases rapidly during the first few iterations and then gradually stabilizes on all eight benchmark datasets. This demonstrates that LCGL can converge rapidly to a stable solution. A similar decreasing trend is observed across datasets with different characteristics, showing that the optimization behavior is consistent in practice. In addition, most of the objective reduction occurs at the early stage, indicating that LCGL requires only a limited number of iterations to reach a near-stationary state.

5 Conclusion

In this paper, we propose Reliable Latent Structure Estimation (ReLaS) for incomplete multi-view semi-supervised learning. ReLaS estimates reliable direct relations through Direct Relation Estimation (DRE), recovers relations rendered unobservable by missing endpoints through Cross-View Relation Recovery (CRR), and integrates both sources through Latent Consensus Graph Learning (LCGL) to obtain a latent consensus graph. Extensive experiments on eight benchmark datasets demonstrate the effectiveness of ReLaS under different view-missing rates, while relation-level analyses further verify its capability to restore relational structures degraded by view missingness. In future work, we will explore more adaptive cross-view relation recovery strategies and extend ReLaS to broader incomplete multi-view learning scenarios.

REFERENCES

Shunshun Bai, Qinghai Zheng, Xiaojin Ren, and Jihua Zhu. Graph-guided imputation-free incomplete multi-view clustering. *Expert Systems with Applications*, 258:125165, 2024. ISSN 0957-4174.

Yiran Cai, Hangjun Che, Baicheng Pan, Man-Fai Leung, Cheng Liu, and Shiping Wen. Projected cross-view learning for unbalanced incomplete multi-view clustering. *Information Fusion*, 105: 102245, 2024. ISSN 1566-2535.

Guoqing Chao, Shiliang Sun, and Jinbo Bi. A survey on multiview clustering. *IEEE Transactions on Artificial Intelligence*, 2(2):146–168, 2021.

Guoqing Chao, Yi Jiang, and Dianhui Chu. Incomplete contrastive multi-view clustering with high-confidence guiding. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI'24/IAAI'24/EAAI'24. AAAI Press, 2024. ISBN 978-1-57735-887-9.

Man-Sheng Chen, Jia-Qi Lin, Xiang-Long Li, Bao-Yu Liu, Chang-Dong Wang, Dong Huang, and Jian-Huang Lai. Representation learning in multi-view clustering: A literature review. *Data Science and Engineering*, 7(3):225–241, 2022.

Wen Chen, Lei Gao, Da Wang, Fang Guo, and Cheng Liang. Geometry-aware variational information maximization for deep incomplete multi-view clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 20289–20297, 2026.

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 *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 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. ISSN 1566-2535.

Xiaojian Ding, Lin Zhao, Xian Li, and Xiaoying Zhu. Incomplete multi-view clustering via hierarchical semantic alignment and cooperative completion. In D. Belgrave, C. Zhang, H. Lin, R. Pascanu, P. Koniusz, M. Ghassemi, and N. Chen (eds.), *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pp. 52175–52195. Curran Associates, Inc., 2025.

Yang Du, Zhen Wang, Jia Li, Rui Ning, and Lei 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.

Hang Gao, Cheng Liu, Hongming Sun, Gaoyang Li, Ying Li, You Zhou, and Wei Du. Incomplete multi-view clustering with cross-view generation via pre-trained transformer. *Pattern Recognition*, 171:112166, 2026. ISSN 0031-3203.

Muhammad Haris, Yusliza Yusoff, Azlan Mohd Zain, Abid Saeed Khattak, and Syed Fawad Hussain. Breaking down multi-view clustering: A comprehensive review of multi-view approaches for complex data structures. *Engineering Applications of Artificial Intelligence*, 132:107857, 2024. ISSN 0952-1976.

Shuo Hu, Bing Tian, Wei Liu, and Yunming Ye. Self-supervised trusted contrastive multi-view clustering with uncertainty refined. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 17305–17313, 2025.

Yanyong Huang, Zongxin Shen, Tianrui Li, and Fengmao Lv. Unified view imputation and feature selection learning for incomplete multi-view data. In Kate Larson (ed.), *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, pp. 4192–4200. International Joint Conferences on Artificial Intelligence Organization, 8 2024. Main Track.

Bingbing Jiang, Chenglong Zhang, Xinyan Liang, Peng Zhou, Jie Yang, Xingyu Wu, Jun-Yi Guan, Weiping Ding, and Weiguo Sheng. Collaborative similarity fusion and consistency recovery for incomplete multi-view clustering. In *AAAI Conference on Artificial Intelligence*, 2025a.

Zhangqi Jiang, Tingjin Luo, Xu Yang, and Xinyan Liang. Adversarial graph fusion for incomplete multi-view semi-supervised learning with tensorial imputation. In D. Belgrave, C. Zhang, H. Lin, R. Pascanu, P. Koniusz, M. Ghassemi, and N. Chen (eds.), *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pp. 51323–51360. Curran Associates, Inc., 2025b.

Jiaqi Jin, Siwei Wang, Zhibin Dong, Xihong Yang, Xinwang Liu, En Zhu, and Kunlun He. Deep incomplete multi-view clustering with distribution dual-consistency recovery guidance. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 1016–1026, 2025.

Zhiyuan Lai, Jia Li, Jinjia Wang, and Shiping Wang. Cooperative graph transformer with structural consensus for multi-view learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 22751–22759, 2026.

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 (IJCAI-23)*, pp. 3911–3919. International Joint Conferences on Artificial Intelligence Organization, 2023.

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(4):4447–4461, 2023.

Qian Liu, Huibing Wang, Jinjia Peng, Yawei Chen, Mingze Yao, Xianping Fu, and Yang Wang. Consensus-guided incomplete multi-view clustering via cross-view affinities learning. In James Kwok (ed.), *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, pp. 5761–5769. International Joint Conferences on Artificial Intelligence Organization, 8 2025. Main Track.

Suyuan Liu, Siwei Wang, Ke Liang, Junpu Zhang, Zhibin Dong, Tianrui Liu, En Zhu, Kunlun He, and Xinwang Liu. Alleviate anchor-shift: Explore blind spots with cross-view reconstruction for incomplete multi-view clustering. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 87509–87531. Curran Associates, Inc., 2024a.

Suyuan Liu, Junpu Zhang, Yi Wen, Xihong Yang, Siwei Wang, Yi Zhang, En Zhu, Chang Tang, Long Zhao, and Xinwang Liu. Sample-level cross-view similarity learning for incomplete multi-view clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 14017–14025, 2024b.

Lilei Sun, Wai Keung Wong, Yusen Fu, Jie Wen, Mu Li, Yuwu Lu, and Lunke Fei. Dual structure-aware consensus graph learning for incomplete multi-view clustering. *Pattern Recognition*, 165: 111582, 2025. ISSN 0031-3203.

Jingjing Tang, Qingqing Yi, Saiji Fu, and Yingjie Tian. Incomplete multi-view learning: Review, analysis, and prospects. *Applied Soft Computing*, 153:111278, 2024. ISSN 1568-4946.

Haiyue Wang, Quan Wang, Qiguang Miao, and Xiaoke Ma. Joint learning of data recovering and graph contrastive denoising for incomplete multi-view clustering. *Information Fusion*, 104:102155, 2024. ISSN 1566-2535.

Jiayuan Wang, Jie Lian, Yongquan Shi, Jielong Lu, and Shiping Wang. Towards multi-order homophily-aware propagation in multi-view learning. *Neural Networks*, 205:109381, 2027.

Shiping Wang, Yueyang Pi, Yang Huang, Fuhai Chen, and Le Zhang. Multi-channel equilibrium graph neural network for multi-view semi-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(10):9375–9382, 2025.

Xinxin Wang, Yongshan Zhang, Xiaochen Yuan, and Yicong Zhou. Cross-view anchor graph learning and factorization for incomplete multi-view clustering. In *AAAI Conference on Artificial Intelligence*, 2026a.

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 C. Vondrick, B. Hariharan, C. Raffel, L. Pinto, D. Yang, and A. Faust (eds.), *International Conference on Learning Representations*, volume 2026, pp. 79726–79740, 2026b.

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 (CVPR)*, pp. 15712–15721, June 2023.

Benyu Wu, Wei Du, Jun Wang, and Guoxian Yu. Imputation-free incomplete multi-view clustering via knowledge distillation. In James Kwok (ed.), *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, pp. 6570–6578. International Joint Conferences on Artificial Intelligence Organization, 8 2025a. Main Track.

Minghong Wu, Jihua Zhu, Wenbiao Yan, Bin Chen, and Qinghai Zheng. Relationship completion for incomplete multi-view clustering. *Neural Networks*, 191:107791, 2025b. ISSN 0893-6080.

Yanwanyu Xi, Xiao Zheng, Chang Tang, Xingchen Hu, Yuanyuan Liu, Jun-Jie Huang, and Xinwang Liu. Lrgr: Self-supervised incomplete multi-view clustering via local refinement and global realignment. In James Kwok (ed.), *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, pp. 6624–6632. International Joint Conferences on Artificial Intelligence Organization, 8 2025. Main Track.

Lei Xing, Badong Chen, Changyuan Yu, and Jing Qin. Robust tensor ring-based graph completion for incomplete multi-view clustering. *Information Fusion*, 111:102501, 2024. ISSN 1566-2535.

Chang Xu, Dacheng Tao, and Chao Xu. A survey on multi-view learning. *arXiv preprint arXiv:1304.5634*, 2013.

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

Yang Yang, De-Chuan Zhan, Xiang-Rong Sheng, and Yuan Jiang. Semi-supervised multi-modal learning with incomplete modalities. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI-18)*, pp. 2998–3004. International Joint Conferences on Artificial Intelligence Organization, 2018.

Zhiwen Yu, Ziyang Dong, Chenchen Yu, Kaixiang Yang, Ziwei Fan, and C. L. Philip Chen. A review on multi-view learning. *Frontiers of Computer Science*, 19(7):197334, 2025.

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*, volume 39, pp. 22650–22658, 2025.

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 *International Joint Conference on Artificial Intelligence*, 2025.

Pengfei Zhu, Xinjie Yao, Yu Wang, Meng Cao, Binyuan Hui, Shuai Zhao, and Qinghua Hu. Latent heterogeneous graph network for incomplete multi-view learning. *IEEE Transactions on Multimedia*, 25:3033–3045, 2023.

Wenzhang Zhuge, Tingjin Luo, Ruidong Fan, Hong Tao, Chenping Hou, and Dongyun Yi. Absent multiview semisupervised classification. *IEEE Transactions on Cybernetics*, 54(3):1708–1721, 2024.

THE USE OF LARGE LANGUAGE MODELS (LLMS)

In this research, large language models (LLMs) were employed exclusively to improve the clarity and grammatical accuracy of the manuscript. Their use was limited to refining sentence structure and correcting syntax to enhance readability and professionalism. At no stage did these tools influence the scientific content, methodology, or results. All core ideas and analyses presented are the original work of the authors, with no LLM-generated content contributing to the intellectual substance of the paper.

A ADDITIONAL MATHEMATICAL DETAILS

This section provides additional derivations and mathematical properties supporting the formulations introduced in the main text.

A.1 POSTERIOR EVIDENCE MODEL

For a direct relation e, conditioning on its O_e observation opportunities gives

$$p(C_e \mid O_e, \theta_e) \propto \theta_e^{C_e}(1-\theta_e)^{O_e - C_e}. \tag{18}$$

Combining this likelihood with the non-informative prior $\theta_e \sim \mathrm{Beta}(1,1)$ yields

$$\begin{aligned} p(\theta_e \mid C_e, O_e) &\propto \theta_e^{C_e}(1-\theta_e)^{O_e - C_e} \\ &= \theta_e^{(C_e+1)-1}(1-\theta_e)^{(O_e - C_e+1)-1}. \end{aligned} \tag{19}$$

Therefore,

$$\theta_e \mid C_e, O_e \sim \mathrm{Beta}\left(C_e + 1,\, O_e - C_e + 1\right), \tag{20}$$

recovering Equation 3. The corresponding posterior variance is

$$\nu_e = \frac{(C_e+1)(O_e - C_e + 1)}{(O_e+2)^2(O_e+3)}, \tag{21}$$

which gives Equation 6. These quantities are interpreted under the working evidence model adopted by DRE.

A.2 PROPERTIES OF CONSERVATIVE CONFIRMATION RELIABILITY

For compactness, define

$$R_\delta(C, O) = Q_\delta\left[\mathrm{Beta}(C+1, O-C+1)\right]. \tag{22}$$

Positive confirmation. For fixed O, the posterior density ratio associated with $(C+1, O)$ and (C, O) satisfies

$$\frac{p(\theta \mid C+1, O)}{p(\theta \mid C, O)} \propto \frac{\theta}{1-\theta}, \tag{23}$$

which is increasing on $(0,1)$. Hence the posterior shifts toward larger confirmation probabilities and

$$R_\delta(C+1, O) \geq R_\delta(C, O). \tag{24}$$

Observable disagreement. For fixed C, adding one observation opportunity without an additional confirmation gives

$$\frac{p(\theta \mid C, O+1)}{p(\theta \mid C, O)} \propto 1-\theta, \tag{25}$$

which is decreasing on $(0,1)$. Consequently,

$$R_\delta(C, O+1) \leq R_\delta(C, O). \tag{26}$$

Thus, positive confirmation increases conservative reliability, whereas an observable disagreement decreases it.

Equal confirmation ratio with different evidence amounts. When every observation opportunity confirms the relation, $C = O$, the posterior reduces to

$$\theta \mid O, O \sim \mathrm{Beta}(O+1, 1). \tag{27}$$

Its cumulative distribution function is $F(\theta) = \theta^{O+1}$, giving

$$R_\delta(O, O) = \delta^{1/(O+1)}. \tag{28}$$

Since $\delta \in (0, 1)$, this quantity increases with O. For example,

$$R_\delta(1, 1) = \delta^{1/2} < \delta^{1/5} = R_\delta(4, 4), \tag{29}$$

although both cases have $C/O = 1$. Conservative reliability therefore depends on both the confirmation ratio and the amount of available evidence.

A.3 Robust Precision Normalization

Equation 7 provides the inverse posterior variance underlying the direct-relation precision. For the final precision used by LCGL, define the raw score

$$\pi_e^{\mathrm{raw}} = \nu_e^{-1}. \tag{30}$$

Let $P_{5\%}$ and $P_{95\%}$ denote the corresponding percentiles of $\{\pi_f^{\mathrm{raw}} : f \in \mathcal{E}_D\}$. Extreme values are first winsorized as

$$\widetilde{\pi}_e = \mathrm{clip}\left(\pi_e^{\mathrm{raw}}; P_{5\%}, P_{95\%}\right). \tag{31}$$

The resulting precision is then normalized by its median:

$$\pi_e = \frac{\widetilde{\pi}_e}{\mathrm{median}_{f \in \mathcal{E}_D} \widetilde{\pi}_f}. \tag{32}$$

Consequently,

$$\mathrm{median}_{e \in \mathcal{E}_D} \pi_e = 1. \tag{33}$$

The transformation preserves the ordering of the winsorized inverse-variance scores while placing their typical scale at one.

A.4 Relation-Profile Regularization

Let

$$d_i^Q = \sum_j Q_{ij}, \qquad D_Q = \mathrm{diag}(d_1^Q, \ldots, d_n^Q), \tag{34}$$

so that

$$\mathcal{L}_Q = I - D_Q^{-1/2} Q D_Q^{-1/2}. \tag{35}$$

The structural term in Equation 16 satisfies

$$\begin{aligned} \mathrm{tr}(S^\top \mathcal{L}_Q S) &= \mathrm{tr}(S^\top S) - \mathrm{tr}\left(S^\top D_Q^{-1/2} Q D_Q^{-1/2} S\right) \\ &= \sum_i \|S_{i,:}\|_2^2 - \sum_{i,j} \frac{Q_{ij}}{\sqrt{d_i^Q d_j^Q}} \langle S_{i,:}, S_{j,:} \rangle . \end{aligned} \tag{36}$$

Using $\sum_j Q_{ij} = d_i^Q$ gives

$$\begin{aligned} &\frac{1}{2} \sum_{i,j} Q_{ij} \left\| \frac{S_{i,:}}{\sqrt{d_i^Q}} - \frac{S_{j,:}}{\sqrt{d_j^Q}} \right\|_2^2 \\ &\quad = \sum_i \|S_{i,:}\|_2^2 - \sum_{i,j} \frac{Q_{ij}}{\sqrt{d_i^Q d_j^Q}} \langle S_{i,:}, S_{j,:} \rangle . \end{aligned} \tag{37}$$

Therefore,

$$\mathrm{tr}(S^\top \mathcal{L}_Q S) = \frac{1}{2} \sum_{i,j} Q_{ij} \left\| \frac{S_{i,:}}{\sqrt{d_i^Q}} - \frac{S_{j,:}}{\sqrt{d_j^Q}} \right\|_2^2, \tag{38}$$

which verifies Equation 17. The structural term therefore promotes consistency between entire relation profiles of samples connected by direct evidence.

A.5 Structural-Curvature Scale

The structural coefficient in Equation 16 balances the typical direct-relation fidelity curvature with the corresponding local curvature of the normalized Laplacian:

$$\lambda = \frac{\operatorname{median}_{e\in\mathcal{E}_D} \pi_e}{\operatorname{median}_{e=(i,j)\in\mathcal{E}_D} \left[(\mathcal{L}_Q)_{ii} + (\mathcal{L}_Q)_{jj}\right]}. \tag{39}$$

From Equation 33,

$$\operatorname{median}_{e\in\mathcal{E}_D} \pi_e = 1. \tag{40}$$

Every endpoint of a direct relation has positive degree in Q, and therefore the normalized Laplacian satisfies

$$(\mathcal{L}_Q)_{ii} = (\mathcal{L}_Q)_{jj} = 1, \qquad e = (i,j) \in \mathcal{E}_D. \tag{41}$$

It follows immediately that

$$\lambda = \frac{1}{2}. \tag{42}$$

A.6 LCGL Optimization and Convexity

Let $\mathbf{s} \in \mathbb{R}^{|\mathcal{E}^+|}$ collect the undirected edge variables of S over $\mathcal{E}^+$. For each $e \in \mathcal{E}^+$, define

$$c_e = \begin{cases} q_e, & e \in \mathcal{E}_D, \\ r_e, & e \in \mathcal{E}_R, \end{cases} \qquad w_e = \begin{cases} \pi_e, & e \in \mathcal{E}_D, \\ \eta_e, & e \in \mathcal{E}_R. \end{cases} \tag{43}$$

Let $\mathbf{c}$ and $\mathbf{w}$ collect these quantities, and define

$$W = \operatorname{diag}(\mathbf{w}). \tag{44}$$

Denote by $\mathcal{M}_{\mathcal{E}^+}(\mathbf{s})$ the linear embedding that places the entries of $\mathbf{s}$ on their corresponding symmetric edges. Equation 16 can then be written as

$$J(\mathbf{s}) = \frac{1}{2}(\mathbf{s}-\mathbf{c})^\top W(\mathbf{s}-\mathbf{c}) + \frac{\lambda}{2} \operatorname{tr}\left[\mathcal{M}_{\mathcal{E}^+}(\mathbf{s})^\top \mathcal{L}_Q \mathcal{M}_{\mathcal{E}^+}(\mathbf{s})\right], \qquad \mathbf{s} \geq 0. \tag{45}$$

For each $e = (i,j) \in \mathcal{E}^+$, the objective gradient with respect to the corresponding undirected edge variable is

$$\frac{\partial J}{\partial s_e} = w_e(s_e - c_e) + \lambda\left[(\mathcal{L}_Q S)_{ij} + (\mathcal{L}_Q S)_{ji}\right]. \tag{46}$$

The resulting problem is solved using L-BFGS-B under the nonnegativity constraints.

For every direct relation, $\pi_e > 0$, while a recovered relation enters $\mathcal{E}_R$ only when $\eta_e > 0$. Hence

$$W \succ 0, \qquad \lambda_{\min}(W) = \min_{e\in\mathcal{E}^+} w_e > 0. \tag{47}$$

The fidelity term is therefore strongly convex in $\mathbf{s}$.

Moreover, $\mathcal{L}_Q \succeq 0$, and for any matrix S,

$$\operatorname{tr}(S^\top \mathcal{L}_Q S) = \sum_{j=1}^{n} S_{:,j}^\top \mathcal{L}_Q S_{:,j} \geq 0. \tag{48}$$

The structural term is thus convex. Since $\mathcal{M}_{\mathcal{E}^+}$ is linear, its composition with this positive semidefinite quadratic form remains convex.

Therefore, Equation 45 is strongly convex with modulus at least

$$\min_{e\in\mathcal{E}^+} w_e > 0. \tag{49}$$

The feasible set $\mathbf{s} \geq 0$ is closed and convex, while the positive-definite fidelity term makes the objective coercive. Consequently, LCGL admits a unique global minimizer $\widehat{S}$.

B Additional Experimental Analysis

B.1 Relation Loss and Recovery Metrics

To quantify the relational degradation caused by view missingness and the recovery capability of CRR, two relation-level metrics are adopted: relation loss ratio and recovery recall.

Given the complete reference relation graph $\mathcal{E}_{\text{full}}$ and the view availability matrix $\mathbf{M}$, the endpoint-induced unobservable relation set is defined as

$$\mathcal{E}_{\text{lost}} = \left\{ (i,j) \in \mathcal{E}_{\text{full}} \mid \sum_{v=1}^{V} M_{iv} M_{jv} = 0 \right\}. \tag{50}$$

The relation loss ratio is then calculated as

$$\text{RelationLossRatio} = \frac{|\mathcal{E}_{\text{lost}}|}{|\mathcal{E}_{\text{full}}|}, \tag{51}$$

which measures the proportion of reference relations that become completely unobservable because no view jointly observes both endpoints.

For relation recovery, only relations that pass reliability calibration and are finally admitted into LCGL are considered. Denoting this recovered relation set as $\mathcal{E}_{\text{rec}}$, the recovery recall is defined as

$$\text{RecoveryRecall} = \frac{|\mathcal{E}_{\text{rec}} \cap \mathcal{E}_{\text{lost}}|}{|\mathcal{E}_{\text{lost}}|}. \tag{52}$$

This metric evaluates how many truly lost relations are successfully recovered and retained by CRR, rather than measuring candidate-level retrieval. The corresponding results are reported in Fig. 4.