Mutually Reinforced Fair Multi-View Clustering

Jiaxin Wang[1] , **Xiaoxu Tan**[2] , **Jieren Cheng**[2,3*] , **Junwei Zheng**[2] , **Wenxuan Tu**[2,3*]
[1]School of Cyberspace Security, Hainan University, Haikou, China
[2]School of Computer Science and Technology, Hainan University, Haikou, China
[3]Hainan Blockchain Technology Engineering Research Center, Haikou, China
{jx_wang2046, xiaoxu_tan}@163.com, {992730, 20243004640, twx}@hainanu.edu.cn

Abstract

Fair multi-view clustering aims to exploit complementary information across views to improve clustering quality, while ensuring unbiased outcomes with respect to sensitive groups. Existing methods typically separate multi-view clustering from fairness optimization into distinct stages. However, such a decoupled design provides limited synergy between the two objectives, which often makes progress on one objective detrimental to the other. To address this challenge, we propose **M**utually **R**einforced **F**air **M**ulti-**V**iew **C**lustering (**MR-FMVC**), which optimizes fairness alignment and clustering iteratively to enable coordinated optimization. Specifically, we first disentangle view representations into common and private components to mitigate fairness-related discrepancies across views. Subsequently, the optimization proceeds by interleaving cross-group matching in the common space with updating cluster centroids in the joint representation space, ultimately assigning matched samples to consistent clusters. Extensive experiments on four fairness datasets demonstrate that MR-FMVC achieves a superior trade-off between clustering performance and fairness.

1 Introduction

Multi-view clustering (MVC) leverages complementary information across multiple views to learn more informative representations and improve clustering quality [Liu *et al.*, 2021; Guan *et al.*, 2025a]. It has been widely adopted in practical decision making systems such as bank credit approval, admissions screening, and medical diagnosis [Liang *et al.*, 2025; Liu *et al.*, 2025]. Nonetheless, sensitive groups may be subject to unequal treatment in real-world datasets, which can bias model predictions [Donini *et al.*, 2018; Hutchinson and Mitchell, 2019; Chouldechova and Roth, 2020]. Consequently, fairness in MVC has become a necessary research direction for equitable automated decision-making.

Fair multi-view clustering (FMVC) focuses on learning cluster assignments that are unbiased with respect to sensitive groups, and encourages consistent assignments across views. For example, in admissions screening, applicants typically submit multi-source materials such as personal statements, letters of recommendation and other supporting materials. It is crucial to ensure that sensitive information (e.g., gender and race) should not adversely affect the final decisions. Existing methods enforce fairness primarily via adversarial learning or explicit fair constraints. FairMVC [Shekhar *et al.*, 2023] is the first to introduce fairness into MVC by explicitly constraining the sensitive group distribution within each cluster to match the global distribution. DFMVC [Zhao *et al.*, 2024] learns consistent multi-view representations with contrastive learning and enforces fairness via a fairness constraint. DFMVC-AKEN [Xu *et al.*, 2025] couples attention with KAN's strong nonlinear representational capacity and further enforces group fairness with the DFMVC fairness constraint. AFMVC [Jiang *et al.*, 2026] gradually removes sensitive information from latent representations via an adversarial mechanism, thereby mitigating the influence of sensitive attributes on downstream decisions. While these methods alleviate unfairness in MVC to some extent, further improving fairness requires consideration of how fair groups are defined and constructed [Chierichetti *et al.*, 2017].

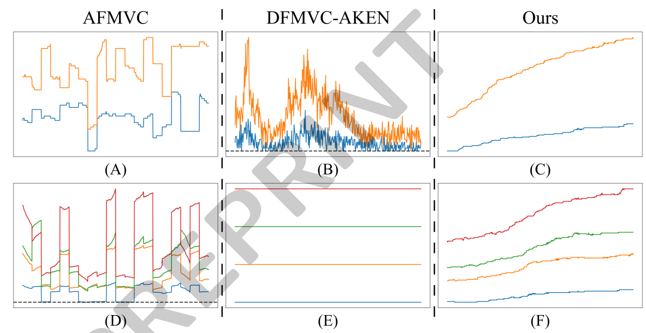

Figure 1: It illustrates the evolution of within-cluster sensitive-group proportions during training on Credit (top, 2 sensitive groups and 2 views) and MFeat (bottom, 4 sensitive groups and 6 views). The dashed line denotes an asymptote approaching zero. As the numbers of sensitive groups and views increase, existing methods exhibit pronounced uncertainty and even converge to shortcut solutions (Fig. E), whereas our method remains stable across settings.

Fair groups refer to basic units constructed by proportion-

*Corresponding author.

ally sampling according to the global distribution of sensitive groups. Performing clustering over such units can attenuate the influence of sensitive attributes at the assignment level. However, in practical scenarios with an increasing number of views or sensitive groups, substantial distributional discrepancies across views can trigger fairness-related discrepancies, such that the same sample may correspond to different fair group constructions under different views. Despite recent progress in FMVC, decoupling clustering from fairness objectives incurs inherent limitations, which are amplified in practical. (i) Most approaches conduct shallow view fusion after feature encoding, where sensitive components from different views become entangled, leading to fairness discrepancies and information redundancy. This design restricts effective exploitation of multi-view complementarity and makes it difficult to learn high-quality fair representations. (ii) Existing methods enforce fairness by aggressively aligning within-cluster sensitive-group distributions with the global distribution. However, fairness optimization is typically imposed as an external objective that is largely decoupled from clustering optimization. As a result, the two objectives pull against each other during optimization, blurring the optimization trajectory and causing mutual interference. As illustrated in Fig. 1, the decoupled objective design leads to an ambiguous optimization trajectory, and this ambiguity is further exacerbated in practical scenarios. Therefore, it is necessary to develop a learning paradigm that (i) constructs a high-quality consensus representation and (ii) coordinates the joint optimization of clustering and fairness constraints in that space.

To overcome the above limitations, we propose MR-FMVC, which builds upon synergistic optimization. Within each view, it disentangles private factors related to sensitive groups from common factors that are consistent across views. To avoid cross-view fairness-related discrepancies, MR-FMVC performs cross-group sensitive matching exclusively on the common representations, while learning the clustering structure in a consensus space formed by the joint shared–private representation. During end to end training, we first learn robust representations via a multi-view variational encoder, and then alternately iterate between sensitive matching in the common space and centroid updates for consensus clustering, enabling fairness alignment and clustering-structure learning to converge in a coordinated manner. Through this joint optimization strategy, MR-FMVC improves fairness while preserving cluster separability, and substantially enhances training stability. Our main contributions are summarized as follows:

- We propose a fairness-oriented bifactor disentanglement strategy to learn representations disentangled from sensitive attributes. By explicitly modeling cross-view common structure and private information, the strategy prevents fairness-related representations discrepancies.
- We propose a match-calibrated fair clustering that performs sensitive matching exclusively in the common space and updates cluster centroids in the consensus space. By iteratively optimizing matching and clustering, the framework alleviates entangled optimal transport and enables coordinated optimization.
- Extensive experiments on four datasets with sensitive attributes demonstrate that MR-FMVC achieves a superior trade-off between clustering quality and fairness.

2 Related Work

2.1 Multi-view Clustering

Multi-view clustering (MVC) was developed as an extension of conventional single-view clustering [Tu *et al.*, 2021; Tu *et al.*, 2024a; Wang *et al.*, 2026], where each sample is described by multiple heterogeneous views. Bickel et al. [Bickel and Scheffer, 2004] first formalized the basic setting of MVC, emphasizing both complementarity and consistency across views. Building on this foundation, a variety of algorithmic families have been developed. Co-regularized spectral clustering [Kumar *et al.*, 2011] introduces cross-view consistency constraints over spectral embeddings. Matrix factorization methods [Gao *et al.*, 2013] achieve view fusion by sharing representations in a low-dimensional latent space. Multiple kernel learning [Liu *et al.*, 2016; Liu *et al.*, 2020] treats the similarity matrix of each view as a kernel and performs weighted fusion. With the advent of deep learning, increasing attention has been paid to learning nonlinear representations and aligning views [Andrew *et al.*, 2013; Wang *et al.*, 2015]. In addition, several surveys [Fang *et al.*, 2023; Liu *et al.*, 2025] have systematically summarized the methodological landscape, application challenges, and future directions of MVC, contributing to a more structured understanding of the field. Despite the pioneering advances of existing MVC methods [Wang *et al.*, 2024; Guan *et al.*, 2025b] in view fusion, their optimization objectives are primarily designed to promote cluster discriminability. As a result, fairness with respect to sensitive attributes is often left unmodeled, which can induce systematically biased clustering outcomes across sensitive groups.

2.2 Fair Clustering

Fair clustering was first formalized by [Chierichetti *et al.*, 2017] via the notion of fairlets, which ensure balanced sensitive group proportions within each cluster. Subsequent work has extended this framework with more flexible fairness-aware optimization objectives [Bera *et al.*, 2019; Huang *et al.*, 2019]. With the rise of deep learning, deep fair clustering methods have gradually emerged, achieving improved performance on high-dimensional data by incorporating mechanisms such as adversarial strategy [Li *et al.*, 2020; Zhang and Davidson, 2021] and latent-variable disentanglement [Zeng *et al.*, 2023]. Meanwhile, driven by the widespread availability of multi-source data, fairness in MVC has attracted increasing attention [Shekhar *et al.*, 2023], motivating research on ensuring fair clustering outcomes from consensus representations learned over heterogeneous sources. Although existing methods have improved fairness in MVC [Shekhar *et al.*, 2023; Zhao *et al.*, 2024; Xu *et al.*, 2025; Jiang *et al.*, 2026], they typically decouple clustering from fairness optimization into separate stages, resulting in limited coordination between the two objectives. Moreover, as the number of sensitive groups or views increases, they often become unstable, limiting scalability to practical scenarios.

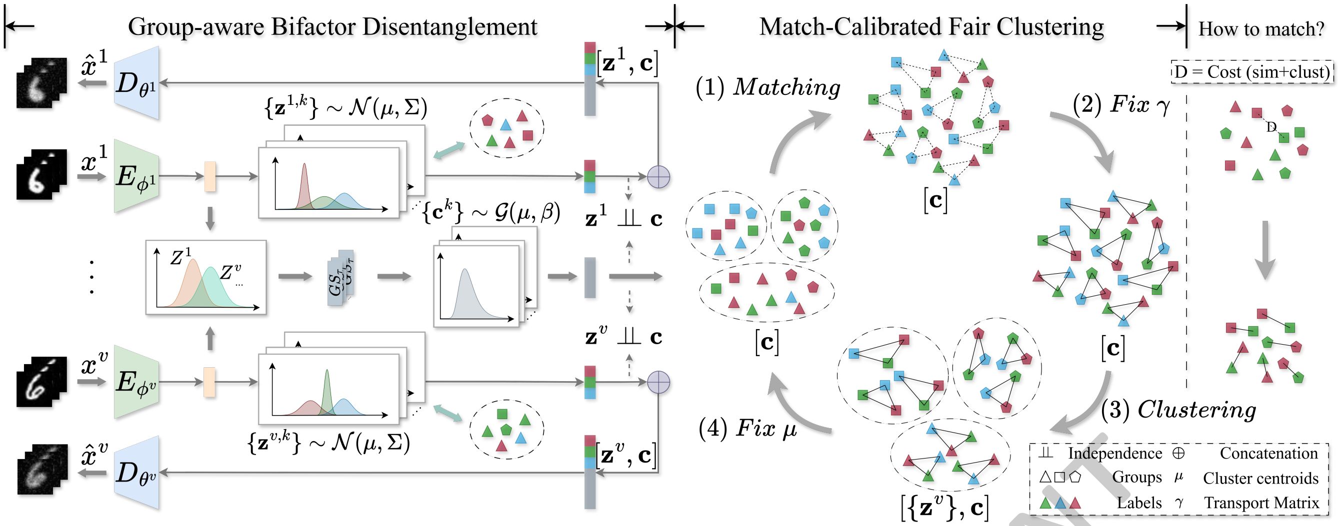

Figure 2: The overall architecture of MR-FMVC. It first learns disentangled multi-view representations with a multi-view variational autoencoder, yielding view-specific factors and a cross-view shared factor, and promotes their independence via a mutual-information minimization regularizer. It then performs fairness-aware clustering on the shared representation via cross-group sensitive matching. We optimize the objective with a coupled alternating scheme that interleaves updates of (i) the matching distribution and (ii) the consensus cluster centroids. When establishing the matching relations, we jointly consider pairwise sample similarity and cluster-level structural information.

3 Method

Notations & Preliminaries Let $\mathcal{X} = \{\mathbf{X}^v\}_{v=1}^{V}$ be a multi-view dataset with $\mathbf{X}^v \in \mathbb{R}^{N\times d_v}$. Each sample i has a sensitive-group label $x_i \in \{0,\dots,G-1\}$, inducing $\mathcal{G}_g = \{i \mid x_i = g\}$ and global proportions $p_g = |\mathcal{G}_g|/N$, and we denote $\mathbf{p} = [p_0,\dots,p_{G-1}]^\top$. A clustering into K clusters is encoded by $\mathbf{z} \in \{1,\dots,K\}^N$, where $\mathcal{C}_k = \{i \mid z_i = k\}$ and $n_k = |\mathcal{C}_k|$. The within-cluster group distribution is $p_{kg} = |\mathcal{C}_k \cap \mathcal{G}_g|/n_k$, and we denote $\mathbf{p}_k = [p_{k0},\dots,p_{k,G-1}]^\top$. Fair clustering requires $\mathbf{p}_k = \mathbf{p}$ for all k, while preserving the inherent clustering structure of the multi-view data.

Existing FMVC methods predominantly rely on adversarial training or hard constraints to enforce fair distributions. However, these mechanisms often induce pronounced multi-objective conflicts. As the numbers of views and sensitive groups increase, sensitive-group distribution alignment becomes increasingly challenging, undermining scalability deployment. To address these challenges, we draw inspiration from the principle in [Kim *et al.*, 2025] and extend it to multi-group matching. The key idea is to learn cross-group sample correspondences that both satisfy the fairness constraints and preserve the clustering structure, and then perform clustering in the induced aligned representation space. Accordingly, we optimize the matching by jointly accounting for pairwise sample similarity and its impact on the clustering objective:

$$C_{\text{match}} = C_{\text{sim}}(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{G-1}) + C_{\text{clust}}(\mathcal{X}, \boldsymbol{\mu}). \quad (1)$$

In multi-view settings, substantial heterogeneity across views makes it difficult to learn matching relations that are consistent across views. This motivates the need for high-quality consensus representations that preserve discriminability while ensuring cross-view matching consistency.

3.1 Group-aware Bifactor Disentanglement

Existing FMVC methods typically perform shallow view fusion after encoding. However, such fusion largely operates as a feature-level combination and therefore cannot explicitly handle pervasive information redundancy and structural disagreement across views. Moreover, sensitive attributes can be implicitly encoded in the latent space and propagated through fusion, causing the final representation to violate fairness. As the number of views increases, redundancy and disagreement accumulate, resulting in further degradation of fairness.

Performing cross-group matching on top of such fusion further amplifies the aforementioned issues, giving rise to entangled optimal transport. Specifically, different views may induce inconsistent matching relations for the same samples, which leads to mutually conflicting transport couplings. Such conflicts compromise global alignment and distort cluster structure. As the number of views grows, the accumulated inconsistencies degrade fairness and clustering performance.

To mitigate these issues, we adopt a view disentanglement scheme [Xu *et al.*, 2021] that factorizes each view representation into a cross-view-consistent shared component and a view-specific private component. We assume independence between the two components, i.e., $p(\mathbf{z}^v, \mathbf{c}) = p(\mathbf{z}^v)p(\mathbf{c})$ and $q(\mathbf{z}^v, \mathbf{c} \mid \mathbf{x}^v) = q(\mathbf{z}^v \mid \mathbf{x}^v)q(\mathbf{c} \mid \mathbf{x}^v)$, and derive Eq. (2) from the β-VAE objective [Higgins *et al.*, 2017]:

$$\mathcal{L} = \sum_{v=1}^{V} \Big[\mathbb{E}_{q(\mathbf{z}^v, \mathbf{c} \mid \{\mathbf{x}^v\})} \log p(\mathbf{x}^v \mid \mathbf{z}^v, \mathbf{c}) - D_{\mathrm{KL}}(q(\mathbf{z}^v \mid \mathbf{x}^v) \,\|\, p(\mathbf{z}^v)) - D_{\mathrm{KL}}(q(\mathbf{c} \mid \{\mathbf{x}^v\}) \,\|\, p(\mathbf{c})) \Big], \quad (2)$$

where $q_\phi(\mathbf{c} \mid \{\mathbf{x}^v\}) = \prod_{k=1}^{K} q_\phi(c_k \mid \{\mathbf{x}^v\})$ denotes the variational posterior of the shared latent variable $\mathbf{c}$ inferred from all views, and $q_{\phi^v}\left(\{\mathbf{z}^{v,k}\}_{k=1}^K \mid \mathbf{x}^v\right) = \prod_{k=1}^K q_{\phi^{v,k}}\left(\mathbf{z}^{v,k} \mid \mathbf{x}^v\right)$ is the factorized variational posterior of the view-private latent variables for view v. Each factor $q_{\phi^{v,k}}\left(\mathbf{z}^{v,k} \mid \mathbf{x}^v\right)$ is modeled as a Gaussian distribution and is sampled via the reparameterization trick in Eq. (3), while $\mathbf{c}$ is sampled from the Gumbel–Softmax distribution parameterized by logits $\mathbf{s}$ as in Eq. (4).

Concretely, given multi-view observations $\{\mathbf{x}^v\}_{v=1}^V$, we parameterize a shared latent variable $\mathbf{c}$ and view-private latents $\mathbf{z}^v$ via variational inference networks. For each view v, a view-specific encoder produces K Gaussian factors $\{\mathbf{z}^{v,k}\}_{k=1}^K$, which are concatenated to form the private representation $\mathbf{z}^v = \mathrm{concat}(\mathbf{z}^{v,1}, \ldots, \mathbf{z}^{v,K})$:

$$\mathbf{z}^{v,k} = \boldsymbol{\mu}^{v,k}(\mathbf{x}^v) + \boldsymbol{\sigma}^{v,k}(\mathbf{x}^v) \odot \boldsymbol{\varepsilon}^{v,k}, \tag{3}$$

where $\boldsymbol{\varepsilon}^{v,k} \sim \mathcal{N}(0, I)$, and $\boldsymbol{\mu}^{v,k}$ and $\boldsymbol{\sigma}^{v,k}$ are output by neural networks given the v-th view embedding.

The shared variable $\mathbf{c}$ is inferred from all views through a cross-view fusion network that outputs logits $\mathbf{s}$, from which we draw a differentiable sample using the Gumbel–Softmax distribution [Jang *et al.*, 2017]:

$$\mathbf{c} = \frac{\exp\big((\log \mathbf{s} + \mathbf{g})/\tau\big)}{\sum_{i=1}^K \exp\big((\log s_i + g_i)/\tau\big)}, \tag{4}$$

where $g_i \sim Gumbel(0, 1)$ and τ is the temperature parameter to control the relaxation.

The resulting posteriors yield a common embedding $\mathbf{c}$ and view-specific embeddings $\{\mathbf{z}^v\}$. We perform cross-group matching exclusively on $\mathbf{c}$ to align fairness-related representations, while using $\{\mathbf{z}^v\}$ only as auxiliary cues for clustering.

3.2 Match-Calibrated Fair Clustering

To balance distributional consistency and cluster discriminability in cross-group sample matching, we decompose the matching objective into two complementary terms: a pairwise similarity cost and a clustering-effect regularizer. The pairwise term penalizes cross-group discrepancies to promote group-invariant alignment, while incorporating group-aware weighting to balance the contributions of different sensitive groups. The clustering-effect term couples matching with the downstream clustering objective in the aligned space, discouraging matchings that distort the cluster geometry and thereby preserving separability. Optimizing the two terms jointly yields a more stable trade-off between fairness and clustering quality. Formally, the resulting objective is:

$$\mathcal{L}_{\text{match}} = \mathbb{E}_{(\mathbf{X}_0,\ldots,\mathbf{X}_{G-1})\sim\mathbb{Q}}\Bigg[\sum_{s=0}^{G-1} G\,\pi_G \Big\| \sum_{\substack{s'=0\\ s'\neq s}}^{G-1} (\mathbf{X}_s - \mathbf{X}_{s'}) \Big\|^2 + \min_k \big\|\mathbf{T} - \boldsymbol{\mu}_k\big\|^2\Bigg], \tag{5}$$

where $\mathbb{Q}$ denotes the joint sampling distribution over the G sensitive groups, $(\mathbf{X}_0, \ldots, \mathbf{X}_{G-1}) \sim \mathbb{Q}$ is a tuple of group-wise representations in the common space, and $\pi_s = \frac{n_s}{n}$ is the group weight. We set $\pi_G = \prod_{s=0}^{G-1} \pi_s$ to weight multi-group matching. The first term penalizes cross-group discrepancies by aggregating pairwise differences between $\mathbf{X}_s$ and $\mathbf{X}_{s'}$, encouraging group-invariant alignment, while the second term $\min_k \|\mathbf{T} - \boldsymbol{\mu}_k\|^2$ preserves clustering structure, where $\mathbf{T} = \sum_{s=0}^{G-1} \pi_s \mathbf{X}_s$ is the weighted barycenter and $\{\boldsymbol{\mu}_k\}_{k=1}^K$ are cluster centroids.

After performing sensitive-group matching in the shared representation space, we construct a transport matrix by jointly considering the pairwise cost and the clustering cost, and use it as a structural prior to guide subsequent clustering updates. We formalize this objective as a Kantorovich optimal transport (OT) problem [Kantorovich, 1960] and solve it efficiently via the Sinkhorn algorithm [Cuturi, 2013]:

$$\min_{\gamma} \ \|(\mathbf{C} + \mathbf{D}) \odot \gamma\|, \ \text{ s.t. } \gamma \mathbf{1} = \frac{1}{n_q}\mathbf{1}, \ \gamma^\top \mathbf{1} = \frac{1}{n_p}\mathbf{1}, \tag{6}$$

where $\gamma \in \mathbb{R}_+^{n_q \times n_p}$ denotes the transport plan between two sensitive-group sample sets of sizes n_q and n_p, and the optimal γ is:

$$\gamma^* = \arg\min_{\gamma \geq 0} \ \langle M, \gamma \rangle - \varepsilon H(\gamma), \tag{7}$$

where $\mathbf{M} = \mathbf{C} + \mathbf{D}$ is the ground cost matrix with entries $\mathbf{m}_{i,j} = \mathbf{c}_{i,j} + \mathbf{d}_{i,j}$, $c_{i,j} = \pi_G \|\mathbf{x}_i - \mathbf{x}_j\|^2$ measures cross-group representation discrepancy, and $d_{i,j} = \min_k \|\mathbf{T} - \boldsymbol{\mu}_k\|^2$ encodes the clustering-effect term with centroids $\{\boldsymbol{\mu}_k\}_{k=1}^K$. The entropy term $H(\gamma) = -\sum_{i,j} \gamma_{ij}(\log \gamma_{ij} - 1)$ regularizes the transport plan, and ε controls the regularization strength.

By estimating an OT plan between the empirical distributions of sensitive groups, we establish alignment relations in the common representation space. Meanwhile, we fuse view-specific private representations with the common space and update cluster centroids in the resulting joint representation to capture global cluster structure. This coupling allows clustering to adapt to alignment-induced geometric changes while retaining view-specific discriminative cues, so that clustering and cross-group matching reinforce each other during optimization, yielding more stable and fair outcomes. Accordingly, we perform soft clustering and incorporate the transport-induced pairings as a structural prior over cluster assignments. The soft assignment is defined as follows:

$$r_k(\mathbf{x}) = \frac{\exp\big(-\|\mathbf{x} - \boldsymbol{\mu}_k\|_2^2/\tau\big)}{\sum_{l=1}^K \exp(-\|\mathbf{x} - \boldsymbol{\mu}_l\|_2^2/\tau)}, \tag{8}$$

where $r_k(\mathbf{x})$ is the soft assignment of $\mathbf{x}$ to cluster k, $\boldsymbol{\mu}_k$ is the centroid, and $\tau > 0$ controls assignment sharpness.

The cluster centroids are updated accordingly:

$$\mathcal{L}_{\text{cluster}} = \min_{\boldsymbol{\mu}} \sum_{0 \leq p < q < G} \sum_{i \in \mathcal{I}_p} \sum_{j \in \mathcal{I}_q} \gamma_{ij}^{(p,q)} \sum_{k=1}^K \Big(r_k(\mathbf{c}_i)\, \|\mathbf{c}_i - \boldsymbol{\mu}_k\|_2^2 + r_k(\mathbf{c}_j)\, \|\mathbf{c}_j - \boldsymbol{\mu}_k\|_2^2 \Big) + \min_{\boldsymbol{\mu}'} \sum_{v=1}^V \sum_{i=1}^N \sum_{k=1}^K r'_k(\mathbf{z}_i^v)\, \|\mathbf{z}_i^v - \boldsymbol{\mu}'_k\|_2^2, \tag{9}$$

where $\mathcal{I}_p$ denotes the index set of samples in the p-th sensitive group, and $\gamma_{ij}^{(p,q)}$ is the transport weight between sample $i \in \mathcal{I}_p$ and sample $j \in \mathcal{I}_q$. $\mathbf{c}_i$ is the common representation of sample i, $\mathbf{z}_i^v$ is the joint representation for view v , and $\boldsymbol{\mu} = \{\boldsymbol{\mu}_k\}_{k=1}^K$ and $\boldsymbol{\mu}' = \{\boldsymbol{\mu}'_k\}_{k=1}^K$ are the centroids in the common space and the joint representation space, respectively.

Building on this design, we adopt an alternating optimization scheme. The updated cluster centroids provide a structural prior and constraints for the next round of matching, while the newly inferred matching relations, in turn, rectify the positions of the centroids in the representation space. Through this bidirectionally coupled iterative process, the model progressively achieves fair alignment across sensitive groups while disentangling structural information, ultimately converging to a representation space that is both fair and well-separable. Finally, we concatenate the common representation with the private representations to form a consensus embedding, and perform clustering updates in this space. The aligned sample pairings are injected into the assignment step as a structural prior, encouraging paired samples to receive consistent cluster memberships:

$$A_s(\mathbf{x}_s)_k := \mathbb{Q}\left(\arg\min_{k'} \left\|\pi_s \mathbf{x}_s + \textstyle\sum_{s' \neq s} \pi_{s'} \mathbf{X}_{s'} - \boldsymbol{\mu}_{k'}\right\|^2 = k \;\middle|\; \mathbf{X}_s = \mathbf{x}_s\right), \tag{10}$$

where $A_s(\mathbf{x}_s)_k$ denotes the probability that a sample $\mathbf{x}_s$ from sensitive group s is assigned to cluster k under the matching-induced joint distribution $\mathbb{Q}$. Conditioned on $\mathbf{X}_s = \mathbf{x}_s$, we construct a joint aligned representation by aggregating $\mathbf{x}_s$ with its matched counterparts from other groups, i.e., $\pi_s \mathbf{x}_s + \sum_{s' \neq s} \pi_{s'} \mathbf{X}_{s'}$, and assign this joint representation to the nearest centroid. The detailed learning procedure of MR-FMVC is shown in Algorithm 1.

3.3 Computational Complexity

We denote by N, V, K, and G the numbers of data points, views, clusters, and sensitive groups, respectively. Let M be the maximum number of neurons in the autoencoders, Z the dimensionality of the view-specific latent variables, n the number of discrete supports per group used in OT, and T the number of Sinkhorn iterations. In practice, $V, K, Z \ll M$.

In GBD, the Multi-VAE module is optimized by minimizing Eq. (2). At each iteration, generating the priors for the shared latent variable costs $\mathcal{O}(NK)$, while generating the priors for the view-specific latent variables costs $\mathcal{O}(VNZ)$. The dominant cost comes from the forward/backward passes through the V autoencoders, which scale as $\mathcal{O}(VNM^2)$. Therefore, the overall per-iteration complexity of this module is $\mathcal{O}\big(NK + VNZ + VNM^2\big)$, which is linear in N.

In MCFC, direct multi-marginal OT requires optimizing a G-way coupling tensor with $\mathcal{O}(n^G)$ variables and comparable memory cost, which quickly becomes intractable as G grows. We instead adopt a pairwise coupling approximation, decomposing it into $\binom{G}{2}$ bi-marginal OT subproblems, each defined on an $\mathcal{O}(n^2)$ coupling matrix, thereby reducing both time and memory overhead. Each bi-marginal OT problem is solved via Sinkhorn iterations in $\mathcal{O}(Tn^2)$ time per pair, leading to an overall matching complexity of $\mathcal{O}\Big(\binom{G}{2}Tn^2\Big)$.

Algorithm 1 MR-FMVC Algorithm

Input: Multi-view dataset $\{\mathbf{x}_i^1, \mathbf{x}_i^2, \ldots, \mathbf{x}_i^V\}_{i=1}^N$; number of clusters K; number of sensitive groups G
Parameter: Regularization hyperparameters (e.g., ε, τ)
Output: Cluster centers $\boldsymbol{\mu} = \{\boldsymbol{\mu}_k\}_{k=1}^K$; assignments $A_s(\mathbf{x}_i)$

// Group-aware Bifactor Disentanglement (GBD)
Optimize multi-view VAE parameters ϕ and θ by minimizing Eq. (2).
Obtain common representation $\mathbf{c}$ and view-private representations $\{\mathbf{z}^v\}_{v=1}^V$ for each view v.
// Match-Calibrated Fair Clustering (MCFC)
Initialize cluster centers $\boldsymbol{\mu} = \{\boldsymbol{\mu}_k\}_{k=1}^K$.
while $\boldsymbol{\mu}$ has not converged **do**
  Compute the transport matrix γ in the common space by solving Eq. (7).
  Update cluster centers $\boldsymbol{\mu}$ in the joint representation space by minimizing Eq. (9).
end while
Assign fair-matched samples by Eq. (10).
return $\boldsymbol{\mu}$ and $\{A_s(\mathbf{x}_i)\}$.

Dataset	Instances	Clusters	Sensitive Attribute	Groups	Views	MaxEnt
Credit	5000	5	Education	2	2	3.3650
Bank	5000	2	Marital	3	2	1.8006
Adult	5000	2	Race	3	3	2.0593
Mfeat	2000	10	Synthetic Value	4	6	12.7985

Table 1: Summary statistics of the four datasets.

4 Experiments

4.1 Experimental Setup

Datasets We evaluate MR-FMVC on four widely used fairness benchmark datasets [Quy *et al.*, 2022] spanning diverse real-world scenarios: Credit, Bank, Adult, and Mfeat. For datasets that are originally single-view, we construct multi-view inputs by applying nonlinear transformations to generate additional views, following [Shekhar *et al.*, 2023]. For datasets without an annotated protected attribute, we synthesize sensitive-group labels by independently sampling a Bernoulli variable for each instance, as in [Zafar *et al.*, 2017]. Dataset statistics are summarized in Table 1.

Baseline Methods To systematically assess the advancement of MR-FMVC, we consider a comprehensive set of baselines spanning three research paradigms, including (i) classical clustering methods, namely K-means [Lloyd, 1982], DEC [Xie *et al.*, 2016], and CC [Li *et al.*, 2021], (ii) representative MVC approaches, including MFLMC [Xu *et al.*, 2022] and DealMVC [Yang *et al.*, 2023], and (iii) recent FMVC algorithms, including F-MVC [Shekhar *et al.*, 2023], DFMVC [Zhao *et al.*, 2024], D-AKEN [Xu *et al.*, 2025], and AFMVC [Jiang *et al.*, 2026]. For conventional clustering baselines that lack explicit multi-view support, we adopt a standard adaptation by concatenating features from all views and applying a single-view clusterer to the fused representation, ensuring consistent inputs and evaluation protocols. Since existing FMVC methods are tailored to binary sensitive attributes, we generalize them to the multi-group setting

Methods	Credit			Bank			Adult			Mfeat		
	ACC(%)	NMI(%)	ENT	ACC(%)	NMI(%)	ENT	ACC(%)	NMI(%)	ENT	ACC(%)	NMI(%)	ENT
K-means	45.08 ± 1.27	21.81 ± 0.54	3.2571 ± 0.062	70.44 ± 1.09	2.62 ± 0.41	1.7919 ± 0.028	68.52 ± 1.34	0.47 ± 0.12	1.9810 ± 0.037	78.40 ± 1.58	80.60 ± 0.95	12.6877 ± 0.064
DEC	45.18 ± 0.73	22.73 ± 1.66	3.2632 ± 0.011	67.60 ± 0.92	4.07 ± 1.37	1.7949 ± 0.011	63.86 ± 0.44	0.63 ± 0.32	1.9828 ± 0.055	79.65 ± 0.89	80.94 ± 1.26	12.7017 ± 0.074
CC	45.50 ± 1.88	20.47 ± 0.29	3.2565 ± 0.076	71.08 ± 1.43	3.49 ± 0.98	1.7919 ± 0.021	69.26 ± 1.59	0.71 ± 0.64	1.9750 ± 0.082	52.85 ± 0.47	54.48 ± 1.35	12.6957 ± 0.008
MFLVC	34.14 ± 0.56	15.12 ± 1.79	3.2604 ± 0.023	58.10 ± 0.37	0.53 ± 0.41	1.7939 ± 0.042	54.80 ± 0.68	2.27 ± 1.17	0.7716 ± 0.093	43.55 ± 0.26	48.27 ± 1.52	7.5672 ± 0.036
DealMVC	46.10 ± 1.41	15.46 ± 0.63	2.2557 ± 0.057	67.36 ± 1.08	0.45 ± 0.34	1.7343 ± 0.097	56.22 ± 1.25	1.91 ± 0.48	0.5983 ± 0.081	58.15 ± 1.56	55.45 ± 0.92	10.1608 ± 0.019
F-MVC	38.36 ± 0.94	19.68 ± 1.37	3.2622 ± 0.084	56.52 ± 0.61	1.88 ± 1.14	1.7991 ± 0.006	54.10 ± 0.35	1.32 ± 1.09	0.8138 ± 0.021	64.50 ± 0.83	67.93 ± 1.64	11.4860 ± 0.049
DFMVC	34.78 ± 1.62	16.15 ± 0.74	3.3631 ± 0.013	55.82 ± 1.91	2.92 ± 0.66	1.7992 ± 0.008	66.32 ± 1.05	14.60 ± 0.39	$\mathbf{2.0592 \pm 0.001}$	10.00 ± 0.00	0.00 ± 0.00	1.27985 ± 0.000
D-AKEN	32.90 ± 0.68	14.73 ± 1.54	$\mathbf{3.3632 \pm 0.016}$	57.48 ± 0.47	0.56 ± 0.22	$\mathbf{1.7995 \pm 0.002}$	67.12 ± 1.35	14.28 ± 0.84	2.0589 ± 0.002	10.00 ± 0.00	0.00 ± 0.00	1.27985 ± 0.000
AFMVC	48.18 ± 1.24	22.94 ± 0.59	3.2734 ± 0.076	74.62 ± 1.15	5.05 ± 0.42	1.7910 ± 0.043	70.90 ± 1.56	0.51 ± 0.21	1.9077 ± 0.034	84.55 ± 0.64	84.56 ± 1.82	12.7165 ± 0.027
Ours	$\mathbf{54.48 \pm 0.83}$	$\mathbf{23.30 \pm 0.61}$	3.3525 ± 0.019	$\mathbf{84.56 \pm 0.73}$	$\mathbf{5.86 \pm 0.92}$	1.7983 ± 0.014	$\mathbf{78.85 \pm 0.26}$	$\mathbf{15.90 \pm 0.68}$	2.0318 ± 0.097	$\mathbf{97.10 \pm 0.52}$	$\mathbf{93.43 \pm 0.41}$	$\mathbf{12.7536 \pm 0.064}$

Table 2: Performance comparison on four fairness datasets. The best and runner-up results are shown in bold and underline, respectively.

by extending the sensitive attribute beyond two groups while strictly preserving their original objective formulations and optimization procedures, thereby maintaining reproducibility, configuration consistency, and credible fairness comparisons.

Metrics We use normalized mutual information (NMI) and clustering accuracy (ACC) to evaluate clustering performance [Tu *et al.*, 2024b]. The conventional Balance (Bal) metric evaluates fairness using the ratio between the largest and smallest sensitive groups. As an extreme-value–based statistic, Bal ignores the distribution of the remaining groups and may fail to reflect global imbalance when many sensitive groups exist. For example, the group ratios $4:4:2:1$ and $4:3:3:1$ yield the same Bal value $4/1$, despite different intermediate-group distributions. To address this limitation, we introduce an Entropy (Ent) metric to measure the global uniformity of sensitive-group distributions. By aggregating all groups, Ent provides a more comprehensive and stable characterization of fairness under multiple sensitive groups.

$$\text{Entropy} = -\sum_{k=1}^{K}\sum_{g=0}^{G-1} P(x_k^g)\,\ln P(x_k^g), \tag{11}$$

where K and G denote the numbers of clusters and sensitive groups, and $P(x_k^g)$ is the empirical proportion of sensitive group g in cluster k.

4.2 Performance Comparison

Table 2 reports the clustering quality and fairness performance of ten methods on four fairness datasets. From these results, some significant observations can be summarized: (1) compared with conventional clustering baselines, MR-FMVC achieves a substantial improvement in Entropy, indicating that it effectively mitigates group bias induced by sensitive attributes and yields fairer cluster assignments; (2) compared with hard-constraint fair clustering approaches (DFMVC and D-AKEN), MR-FMVC attains slightly lower fairness scores yet remains close to perfect fairness, while delivering markedly better clustering performance, suggesting that hard constraints often enforce fairness at the expense of clustering structure, whereas MR-FMVC better preserves the intrinsic data geometry under strong fairness guarantees; (3) compared with the adversarial fair strategy (AFMVC), MR-FMVC adopts a coordinated optimization scheme that simultaneously optimizes clustering quality and fairness, thereby achieving a more reasonable and stable trade-off between the two objectives; and (4) under more challenging data distributions (e.g., Mfeat), several competing methods exhibit clustering degradation or even converge to shortcut solutions, whereas MR-FMVC consistently maintains stable clustering quality and fairness, demonstrating stronger robustness.

Dataset	Credit								
#Groups	2			3			4		
MaxEnt	3.365			5.057			6.2755		
Metrics	ACC(%)	NMI(%)	ENT	ACC(%)	NMI(%)	ENT	ACC(%)	NMI(%)	ENT
F-MVC	38.36	19.68	3.2622	37.56	21.21	4.9466	39.08	13.20	6.0850
DFMVC	34.78	16.15	3.3631	36.68	16.73	**5.0537**	35.96	19.03	6.2662
D-AKEN	32.90	14.73	**3.3632**	37.34	17.85	5.0532	34.02	17.22	**6.2725**
AFMVC	48.18	22.94	3.2734	48.86	21.74	4.9266	47.83	20.81	6.0794
Ours	**54.48**	**23.30**	3.3525	**54.20**	**22.99**	5.0104	**53.88**	**22.30**	6.2339

Table 3: Results on Credit under different group settings.

Dataset	Mfeat								
#Groups	2			3			4		
MaxEnt	6.3651			10.1140			12.7985		
Metrics	ACC(%)	NMI(%)	ENT	ACC(%)	NMI(%)	ENT	ACC(%)	NMI(%)	ENT
F-MVC	64.10	69.62	5.7325	64.15	68.65	9.7369	64.50	67.93	11.4860
DFMVC	10.00	0.00	0.6365	10.00	0.00	1.0114	10.00	0.00	1.27985
D-AKEN	10.00	0.00	0.6365	10.00	0.00	1.0114	10.00	0.00	1.27985
AFMVC	87.55	84.28	6.3082	86.42	84.24	10.0363	84.55	84.56	12.7165
Ours	**98.00**	**95.48**	**6.3484**	**97.30**	**93.92**	**10.0644**	**97.10**	**93.43**	**12.7536**

Table 4: Results on Mfeat under different group settings.

4.3 Different Group Settings

We further vary the number of sensitive groups on the Credit and Mfeat datasets to systematically examine how group cardinality affects model performance. As shown in Tables 3 and 4, as the number of groups increases, existing FMVC methods exhibit optimization instability to different extents. This is mainly because these methods typically adopt a decoupled training paradigm for clustering quality and fairness, which becomes increasingly difficult to optimize jointly under high-dimensional, multi-group settings. Such decoupling can lead to pronounced training oscillations and, in extreme cases, optimization collapse. In contrast, our method alleviates these issues by enforcing coordinated convergence of the two objectives, thereby achieving a more robust trade-off under more challenging configurations.

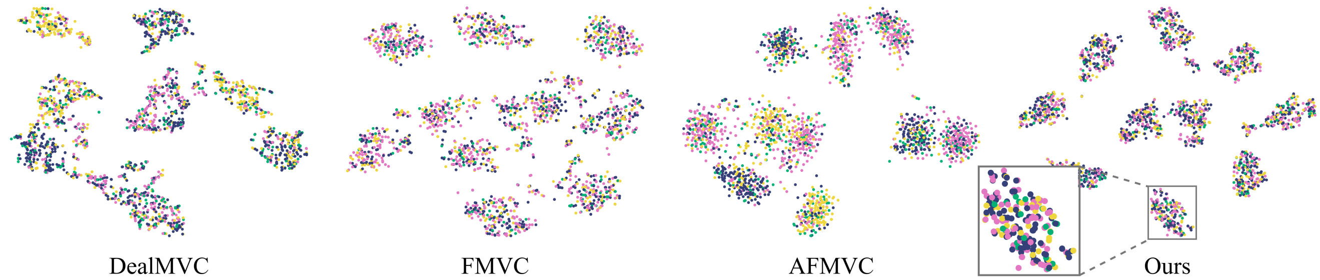

Figure 3: T-SNE Visualization. It corresponds to the results on Mfeat. We use t-SNE to visualize all samples in the learned latent space. MR-FMVC exhibits clearer partitions, more compact clusters, and a more evenly mixed distribution of sensitive groups than competing methods.

MCFC	GBD	Credit			Bank			Adult			Mfeat		
		ACC(%)	NMI(%)	ENT	ACC(%)	NMI(%)	ENT	ACC(%)	NMI(%)	ENT	ACC(%)	NMI(%)	ENT
✓		41.04	19.72	3.2922	74.06	4.19	1.7976	65.45	9.11	1.9916	85.65	80.13	12.7255
	✓	31.08	11.57	3.2965	62.04	0.98	1.7960	59.20	3.47	1.9947	44.95	39.17	12.7444
✓	✓	54.48	23.30	3.3525	84.56	5.86	1.7983	78.85	15.90	2.0318	97.10	93.43	12.7536

Table 5: Ablation study of MR-FMVC. ✓ indicates that the corresponding module (MCFC or GBD) is enabled in each variant.

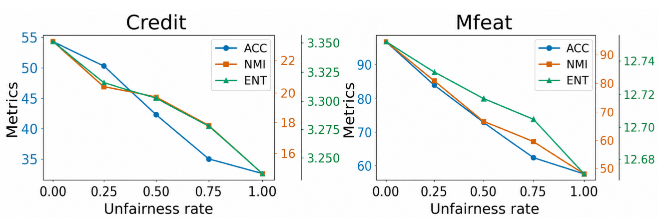

Figure 4: MR-FMVC performance across unfairness rates

4.4 T-SNE Visualization

We visualize the learned representations on MFeat using t-SNE [Maaten and Hinton, 2008]. In Fig. 3, the samples with different colors indicate different sensitive groups predicted by methods. The visualization suggests that existing methods struggle to simultaneously preserve fairness and clustering quality under more challenging settings. In particular, D-AKEN suffers from gradient explosion during training, which leads to optimization failure. In contrast, our method better preserves the clustering structure while ensuring fairness.

4.5 Ablation Study

In this section, we investigate the impact of the MCFC and GBD modules through targeted ablations and analyses. Table 5 reports the comparative results on four datasets. We observe that introducing the MCFC module consistently improves both fairness metrics and clustering-quality metrics, indicating that it helps mitigate sensitive-attribute bias while promoting more discriminative cluster structure. Meanwhile, the GBD module is designed to alleviate discrepancies among multi-view representations and enhance cross-view consistency, which also contributes positively to performance gains. Overall, each component of MR-FMVC effectively improves fairness and clustering performance, and their synergy further strengthens the effectiveness and robustness of MR-FMVC.

4.6 Unfairness Rate

The unfairness rate is a key hyperparameter for practical deployment. Specifically, we rank samples by their matching costs and mask out a subset with the lowest costs, thereby excluding them from the matching procedure. Matching is then performed only on the remaining samples, which yields a controllable relaxation of the fairness constraint. As shown in Fig. 4, increasing the unfairness rate consistently degrades all metrics. This trend indicates that while relaxing the constraint can ease optimization by weakening the matching-induced structural prior, it also reduces fairness and compromises cluster structure. Overall, the unfairness rate provides an interpretable control knob in real-world settings: higher values relax fairness enforcement and reduce matching computation, potentially improving efficiency at the expense of fairness and clustering performance.

5 Conclusion

In this paper, we analyze existing FMVC optimization paradigms and identify a key limitation: clustering and fairness objectives are weakly coupled. To address this issue, we propose MR-FMVC as a new paradigm that iterates between multi-view clustering and fairness alignment, enabling the two objectives to converge in a coordinated manner. Specifically, MR-FMVC incorporates a multi-view sensitive-information disentanglement mechanism and a fair matching-based clustering strategy, which together help maintain a stable balance between fairness and clustering quality in multi-group settings. Extensive experiments on four fairness datasets demonstrate that MR-FMVC effectively mitigates optimization trajectory ambiguity induced by objective decoupling and substantially outperforms competing methods. Future work will extend this framework to more general fairness settings, including intersectional sensitive attributes and dynamically evolving group definitions.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (NSFC) (Grant No. 62562026, 62506102) and the Natural Science Foundation of Hainan University (Grant No. XJ2400009401).

References

[Andrew *et al.*, 2013] Galen Andrew, Raman Arora, Jeff A. Bilmes, and Karen Livescu. Deep canonical correlation analysis. In *Proceedings of the International Conference on Machine Learning*, pages 1247–1255, 2013.

[Bera *et al.*, 2019] Suman Kalyan Bera, Deeparnab Chakrabarty, Nicolas Flores, and Maryam Negahbani. Fair algorithms for clustering. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 4955–4966, 2019.

[Bickel and Scheffer, 2004] Steffen Bickel and Tobias Scheffer. Multi-view clustering. In *Proceedings of the International Conference on Data Mining*, pages 19–26, 2004.

[Chierichetti *et al.*, 2017] Flavio Chierichetti, Ravi Kumar, Silvio Lattanzi, and Sergei Vassilvitskii. Fair clustering through fairlets. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 5029–5037, 2017.

[Chouldechova and Roth, 2020] Alexandra Chouldechova and Aaron Roth. A snapshot of the frontiers of fairness in machine learning. *Communications of the ACM*, 63(5):82–89, 2020.

[Cuturi, 2013] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 2292–2300, 2013.

[Donini *et al.*, 2018] Michele Donini, Luca Oneto, Shai Ben-David, John Shawe-Taylor, and Massimiliano Pontil. Empirical risk minimization under fairness constraints. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 2796–2806, 2018.

[Fang *et al.*, 2023] Uno Fang, Man Li, Jianxin Li, Longxiang Gao, Tao Jia, and Yanchun Zhang. A comprehensive survey on multi-view clustering. *IEEE Transactions on Knowledge and Data Engineering*, 35(12):12350–12368, 2023.

[Gao *et al.*, 2013] Jing Gao, Jiawei Han, Jialu Liu, and Chi Wang. Multi-view clustering via joint nonnegative matrix factorization. In *Proceedings of the International Conference on Data Mining*, pages 252–260, 2013.

[Guan *et al.*, 2025a] Renxiang Guan, Wenxuan Tu, Dayu Hu, Weixuan Liang, Ke Liang, Yaowen Hu, Yue Liu, and Xinwang Liu. Prototype-driven multi-view attribute-missing graph clustering. *IEEE Transactions on Multimedia*, 27:9454–9466, 2025.

[Guan *et al.*, 2025b] Renxiang Guan, Wenxuan Tu, Dayu Hu, Weixuan Liang, Ke Liang, Yaowen Hu, Yue Liu, and Xinwang Liu. Prototype-driven multi-view attribute-missing graph clustering. *IEEE Transactions on Multimedia*, 27:9454–9466, 2025.

[Higgins *et al.*, 2017] Irina Higgins, Loïc Matthey, Arka Pal, Christopher P. Burgess, Xavier Glorot, Matthew M. Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *Proceedings of the International Conference on Learning Representations*, 2017.

[Huang *et al.*, 2019] Lingxiao Huang, Shaofeng H.-C. Jiang, and Nisheeth K. Vishnoi. Coresets for clustering with fairness constraints. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 7587–7598, 2019.

[Hutchinson and Mitchell, 2019] Ben Hutchinson and Margaret Mitchell. 50 years of test (un)fairness: Lessons for machine learning. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 49–58, 2019.

[Jang *et al.*, 2017] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *Proceedings of the International Conference on Learning Representations*, 2017.

[Jiang *et al.*, 2026] Mudi Jiang, Jiahui Zhou, Lianyu Hu, Xinying Liu, Zengyou He, and Zhikui Chen. Adversarial fair multi-view clustering. *Pattern Recognition*, 179:113708, 2026.

[Kantorovich, 1960] Leonid V Kantorovich. Mathematical methods of organizing and planning production. *Management science*, 6(4):366–422, 1960.

[Kim *et al.*, 2025] Kunwoong Kim, Jihu Lee, Sangchul Park, and Yongdai Kim. Fair clustering via alignment. In *Proceedings of the International Conference on Machine Learning*, 2025.

[Kumar *et al.*, 2011] Abhishek Kumar, Piyush Rai, and Hal Daumé III. Co-regularized multi-view spectral clustering. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 1413–1421, 2011.

[Li *et al.*, 2020] Peizhao Li, Han Zhao, and Hongfu Liu. Deep fair clustering for visual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9070–9079, 2020.

[Li *et al.*, 2021] Yunfan Li, Peng Hu, Jerry Zitao Liu, Dezhong Peng, Joey Tianyi Zhou, and Xi Peng. Contrastive clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8547–8555, 2021.

[Liang *et al.*, 2025] Ke Liang, Lingyuan Meng, Hao Li, Jun Wang, Long Lan, Miaomiao Li, Xinwang Liu, and Huaimin Wang. From concrete to abstract: Multi-view clustering on relational knowledge. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(10):9043–9060, 2025.

[Liu *et al.*, 2016] Xinwang Liu, Yong Dou, Jianping Yin, Lei Wang, and En Zhu. Multiple kernel k-means clustering with matrix-induced regularization. In *Proceedings of the*

AAAI Conference on Artificial Intelligence, pages 1888–1894, 2016.

[Liu *et al.*, 2020] Xinwang Liu, Xinzhong Zhu, Miaomiao Li, Lei Wang, En Zhu, Tongliang Liu, Marius Kloft, Dinggang Shen, Jianping Yin, and Wen Gao. Multiple kernel kk-means with incomplete kernels. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(5):1191–1204, 2020.

[Liu *et al.*, 2021] Xinwang Liu, Li Liu, Qing Liao, Siwei Wang, Yi Zhang, Wenxuan Tu, Chang Tang, Jiyuan Liu, and En Zhu. One pass late fusion multi-view clustering. In *Proceedings of the International Conference on Machine Learning*, pages 6850–6859, 2021.

[Liu *et al.*, 2025] Xinwang Liu, Ke Liang, Jun Wang, Suyuan Liu, Xiangke Wang, and Huaimin Wang. Two decades of multi-view clustering: Taxonomy, application, and challenge. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 48(3):3744–3764, 2025.

[Lloyd, 1982] Stuart P. Lloyd. Least squares quantization in PCM. *IEEE transactions on information theory*, 28(2):129–136, 1982.

[Maaten and Hinton, 2008] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11):2579–2605, 2008.

[Quy *et al.*, 2022] Tai Le Quy, Arjun Roy, Vasileios Iosifidis, Wenbin Zhang, and Eirini Ntoutsi. A survey on datasets for fairness-aware machine learning. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 12(3):e1452, 2022.

[Shekhar *et al.*, 2023] Shashi Shekhar, Zhi-Hua Zhou, Yao-Yi Chiang, and Gregor Stiglic. Fairness-aware multi-view clustering. In *Proceedings of the International Conference on Data Mining*, pages 856–864, 2023.

[Tu *et al.*, 2021] Wenxuan Tu, Sihang Zhou, Xinwang Liu, Xifeng Guo, Zhiping Cai, En Zhu, and Jieren Cheng. Deep fusion clustering network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9978–9987, 2021.

[Tu *et al.*, 2024a] Wenxuan Tu, Renxiang Guan, Sihang Zhou, Chuan Ma, Xin Peng, Zhiping Cai, Zhe Liu, Jieren Cheng, and Xinwang Liu. Attribute-missing graph clustering network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 15392–15401, 2024.

[Tu *et al.*, 2024b] Wenxuan Tu, Sihang Zhou, Xinwang Liu, Chunpeng Ge, Zhiping Cai, and Yue Liu. Hierarchically contrastive hard sample mining for graph self-supervised pre-training. *IEEE Transactions on Neural Networks and Learning Systems*, 35(11):16748–16761, 2024.

[Wang *et al.*, 2015] Weiran Wang, Raman Arora, Karen Livescu, and Jeff A. Bilmes. On deep multi-view representation learning. In *Proceedings of the International Conference on Machine Learning*, pages 1083–1092, 2015.

[Wang *et al.*, 2024] Siwei Wang, Xinwang Liu, Suyuan Liu, Wenxuan Tu, and En Zhu. Scalable and structural multi-view graph clustering with adaptive anchor fusion. *IEEE Transactions on Image Processing*, 33:4627–4639, 2024.

[Wang *et al.*, 2026] Jiaxin Wang, Wenxuan Tu, Lingren Wang, Jieren Cheng, and Yue Yang. Anchor-driven nyström for deep graph-level clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 26260–26268, 2026.

[Xie *et al.*, 2016] Junyuan Xie, Ross B. Girshick, and Ali Farhadi. Unsupervised deep embedding for clustering analysis. In *Proceedings of the International Conference on Machine Learning*, pages 478–487, 2016.

[Xu *et al.*, 2021] Jie Xu, Yazhou Ren, Huayi Tang, Xiaorong Pu, Xiaofeng Zhu, Ming Zeng, and Lifang He. Multi-vae: Learning disentangled view-common and view-peculiar visual representations for multi-view clustering. In *Proceedings of the International Conference on Computer Vision*, pages 9214–9223, 2021.

[Xu *et al.*, 2022] Jie Xu, Huayi Tang, Yazhou Ren, Liang Peng, Xiaofeng Zhu, and Lifang He. Multi-level feature learning for contrastive multi-view clustering. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, pages 16030–16039, 2022.

[Xu *et al.*, 2025] Haiming Xu, Qianqian Wang, Boyue Wang, and Quanxue Gao. Deep fair multi-view clustering with attention KAN. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5061–5070, 2025.

[Yang *et al.*, 2023] Xihong Yang, Jiaqi Jin, Siwei Wang, Ke Liang, Yue Liu, Yi Wen, Suyuan Liu, Sihang Zhou, Xinwang Liu, and En Zhu. Dealmvc: Dual contrastive calibration for multi-view clustering. In *Proceedings of the International Conference on Multimedia*, pages 337–346, 2023.

[Zafar *et al.*, 2017] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez-Rodriguez, and Krishna P. Gummadi. Fairness constraints: Mechanisms for fair classification. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, pages 962–970, 2017.

[Zeng *et al.*, 2023] Pengxin Zeng, Yunfan Li, Peng Hu, Dezhong Peng, Jiancheng Lv, and Xi Peng. Deep fair clustering via maximizing and minimizing mutual information: Theory, algorithm and metric. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23986–23995, 2023.

[Zhang and Davidson, 2021] Hongjing Zhang and Ian Davidson. Deep fair discriminative clustering. *CoRR*, abs/2105.14146, 2021.

[Zhao *et al.*, 2024] Bowen Zhao, Qianqian Wang, Zhiqiang Tao, Wei Feng, and Quanxue Gao. DFMVC: deep fair multi-view clustering. In *Proceedings of the International Conference on Multimedia*, pages 8090–8099, 2024.