Fairness-Aware Late-Fusion Multi-View Clustering with Balance and Robustness

Tongzheng Zhao
Shanxi University
Taiyuan, China
202301001148@email.sxu.edu.cn

Yan Chen
Taiyuan University of Technology
Taiyuan, China
chenyan01@tyut.edu.cn

Liang Du*
Shanxi University
Taiyuan, China
duliang@sxu.edu.cn

Abstract

Late-fusion multi-view clustering (LFMVC) offers a scalable solution for analyzing visual and multimodal data by aggregating pre-computed partitions. However, for human-centric visual media applications, existing methods face a critical trade-off: they often lack robustness to noisy views and overlook the crucial requirements of demographic fairness and cluster balance. This paper introduces a unified framework that confronts these three challenges simultaneously. Our approach integrates a multi-level mechanism: (1) a fairness-aware graph filter that purifies base partitions at the representation level, enhancing robustness while mitigating bias; (2) a group-normalized regularizer at the objective level that promotes balanced cluster sizes within each demographic group to induce fairness; and (3) a fairness-guided discretization step that aligns the final assignments with discrete fairness metrics. Extensive experiments on nine benchmarks, including standard visual media datasets, demonstrate that our method achieves state-of-the-art fairness and balance. Crucially, it accomplishes this while maintaining highly competitive clustering accuracy, significantly outperforming specialized fairness baselines and matching or exceeding the performance of state-of-the-art late-fusion methods. The results validate our framework as a practical consensus layer for building robust, equitable, and balanced clustering solutions for visual media analysis. The code for our method is publicly available at https://github.com/Whale-Waves/FLFMVC.

Keywords: Clustering, Fairness, Robustness, Balance

1. Introduction

In the era of large-scale visual media, analyzing heterogeneous data from multiple sources is a fundamental challenge. Consider a content moderation system analyzing a video: it might process visual frames through a CNN for object detection, audio tracks via a speech-to-text model, and user-generated comments using an NLP model. Each source provides a unique "view" of the content. In this context, *late-fusion multi-view clustering* (LFMVC) has emerged as a powerful and scalable paradigm [13]. Instead of centralizing raw, high-dimensional data from diverse sensors or feature extractors, LFMVC intelligently aggregates pre-computed clusterings from each view. This decentralized approach is not only computationally efficient but also respects data privacy and modularity, making it ideal for real-world visual media systems.

Despite its practical advantages, deploying LFMVC in applications involving human-centric visual data—such as personalized content recommendation, demographic analysis from social media images, or facial expression recognition—reveals three critical limitations in existing methods: (*i*) **vulnerability to noise**, where a corrupted view (e.g., from a blurry camera) can degrade the entire consensus; (*ii*) **neglect of demographic fairness**, potentially leading to biased outcomes where clusters disproportionately represent certain groups (e.g., by race or gender); and (*iii*) **disregard for cluster balance**, which is vital for downstream tasks like balanced dataset creation for training generative models or equitable resource allocation. A consensus that is brittle, unfair, or imbalanced undermines the reliability and ethical integrity expected of modern visual media applications.

These issues have largely been addressed in isolation. The LFMVC literature has focused on robustness, primarily through graph-based filtering and alignment techniques that fuse partitions without revisiting raw features [15, 9]. However, these methods are fairness-agnostic and do not enforce balanced partitions. In parallel, the fair clustering community has developed algorithms with strong group-level guarantees [2, 1, 5], but these often assume clean, single-view data and can exacerbate cluster imbalance [4]. A third line on balanced clustering aims to equalize cluster sizes but typically lacks the multi-view robustness needed for complex visual media scenarios.

To date, no LFMVC framework integrates robustness, fairness, and balance within a single principled formulation. This gap forces practitioners to make a difficult choice be-

*Corresponding author.

tween accuracy, fairness, and balance, leading to suboptimal outcomes in domains where all three are essential.

We address this gap with **FLFMVC**, a unified LFMVC framework that jointly optimizes for robustness, fairness, and balance through a *multi-level fairness mechanism*. At the representation level, a fairness-aware graph filter purifies base partitions while embedding fairness priors into the consensus representation. At the objective level, a group-normalized regularizer enforces balanced cluster sizes within each protected group, yielding scale-invariant fairness. During discretization, a fairness-guided refinement step directly optimizes a discrete fairness metric, bridging the gap between continuous relaxation and final assignments. This integration resolves the fairness–accuracy–balance trade-off that constrains existing approaches. Extensive experiments on nine benchmark datasets, including classic visual media collections, demonstrate that FLFMVC achieves state-of-the-art fairness and balance with minimal average accuracy degradation—and in many cases, even improves clustering quality.

2. Related Work

Our work is positioned at the intersection of three key research areas: late-fusion multi-view clustering, fair clustering, and balanced clustering. We review the state-of-the-art in each and highlight the research gap our unified framework aims to fill, particularly within the context of visual media analysis.

2.1. Late Fusion Multi-View Clustering

Late-fusion multi-view clustering (LFMVC) provides a scalable paradigm for consensus learning by combining pre-computed partitions from individual views [9, 11], such as representations from different neural network backbones or data from heterogeneous sensors like cameras and LiDAR. A primary focus is enhancing robustness against noisy views (e.g., images affected by motion blur or poor lighting). Graph-based filtering has emerged as a powerful technique, with methods like sLGm [12] and the highly scalable CFGFLF [3] demonstrating state-of-the-art performance. Other notable approaches include OPLF [9], ALMVC [14], LFLKA [10], and M3LF [7].

Limitations: Despite their effectiveness, existing LFMVC methods are fundamentally fairness-agnostic. They lack mechanisms to prevent partitions that are biased against demographic groups, a critical oversight that limits their applicability in socially sensitive visual media applications, such as facial analysis or automated content moderation.

2.2. Fair Clustering

The field of fair clustering aims to mitigate algorithmic bias by ensuring protected groups are treated equitably, for instance, preventing face recognition systems from performing differently on varying skin tones or genders. Foundational spectral methods like SpFC [6] provide theoretical guarantees for single-view data. More recently, methods like FMSC [8] have begun to extend these concepts to multi-view scenarios.

Limitations: A critical shortcoming of most fair clustering algorithms is their implicit assumption of clean, reliable data. They are not designed to handle the noise and inconsistencies inherent in multi-view visual settings, where views may be corrupted by sensor noise or challenging environmental conditions. Furthermore, they often neglect cluster balance, which can hinder downstream visual applications.

2.3. Balanced and Capacitated Clustering

A third line of research addresses the practical need for clusters of uniform size, which is vital for visual applications like creating balanced mini-batches for training deep neural networks or ensuring user groups in a recommender system are targeted proportionally. Many methods, such as KFC [4], operate on single-view data. More advanced methods like FCE [16] incorporate fairness in an ensemble setting, similar to LFMVC.

Limitations: While these methods excel at controlling cluster sizes, single-view approaches like KFC cannot integrate information from multiple, potentially conflicting visual representations. Although ensemble-based methods like FCE address fairness and balance, they do not incorporate the explicit, theoretically-grounded robustness mechanisms against noisy views that are central to modern LFMVC frameworks and essential for real-world visual data.

2.4. Our Contribution

To the best of our knowledge, the challenges of robustness, fairness, and balance have been addressed largely in isolation. No existing framework provides an integrated solution that combines a state-of-the-art robustness mechanism with explicit fairness and balance regularizers for the LFMVC setting, particularly for visual media analysis. Our proposed FLFMVC bridges this crucial gap by being the first to unify these three objectives within a single, principled optimization framework.

3. The Proposed Method

We propose **FLFMVC**, a late-fusion multi-view clustering framework that targets three practical requirements: robustness to noisy views, demographic fairness, and balanced cluster sizes. Given v pre-computed base partitions $\{\mathbf{H}_r\}_{r=1}^{v}$, each possibly corrupted by noise or bias, FLFMVC learns a consensus clustering $\mathbf{Y}$ and its continuous embedding $\mathbf{F}$. The framework consists of three parts:

(i) a fairness-aware bipartite graph filter with proportional anchor sampling, which reduces sensitivity to weak or noisy connectivity; (ii) a group-normalized regularizer that discourages within-group collapse into a few clusters; and (iii) a fairness-guided discretization step that aligns the final one-hot assignment with a discrete fairness metric.

Let n be the total number of samples and c be the number of clusters. We consider a set of G disjoint protected groups, where group $g \in \{1, \ldots, G\}$ contains n_g samples. Let $\mathcal{I}_g$ denote the index set of samples belonging to group g. The final consensus assignment is represented by $\mathbf{Y} \in \{0,1\}^{n\times c}$, where $Y_{ik} = 1$ iff sample i belongs to cluster k, with the one-hot constraint $\mathbf{Y}\mathbf{1}_c = \mathbf{1}_n$. Consequently, $\mathbf{Y}_g$ denotes the sub-matrix of $\mathbf{Y}$ containing the rows indexed by $\mathcal{I}_g$, i.e., $\mathbf{Y}_g = \mathbf{Y}_{\mathcal{I}_g,:}$, representing the assignments for group g. Similarly, $\mathbf{F} \in \mathbb{R}^{n\times c}$ is the continuous consensus embedding with $\mathbf{F}^\top\mathbf{F} = \mathbf{I}_c$ when optimized, and $\mathbf{F}_g = \mathbf{F}_{\mathcal{I}_g,:}$ is its corresponding sub-matrix for group g. For view r, $\mathbf{W}_r \in \mathbb{R}^{c\times c}$ is an orthogonal alignment matrix ($\mathbf{W}_r^\top\mathbf{W}_r = \mathbf{I}_c$), β_r and γ_r are non-negative weights for partition fusion and graph filtering, $\mathbf{P}_r$ is the normalized bipartite graph, and $\mathbf{G}_r$ its filtered form.

3.1. Fairness-Aware Graph Filtering

We adopt the scalable bipartite graph filtering scheme [3] and modify the anchor selection to avoid over-representing any protected group. Given m anchors, we sample m_g anchors from group g proportionally to its size, i.e., $m_g \approx m \cdot \frac{n_g}{n}$. This keeps cross-group links in the bipartite graph and reduces the risk that a minority group becomes weakly connected.

Bipartite graph construction. For view r, let $\mathbf{B}_r = [b_{ij}^{(r)}] \in \mathbb{R}_{\geq 0}^{n\times m}$ be the sample-to-anchor affinity matrix [3], where $b_{ij}^{(r)}$ is the (nonnegative) similarity between sample i and anchor j (row-sparse via k-NN anchors). We select m anchors by group-wise k-means with $m_g \approx m \cdot \frac{n_g}{n}$ anchors for group g, and compute $\mathbf{B}_r$ by connecting each sample to its k nearest anchors with adaptive weights. Define the degree matrices $\mathbf{D}_n^{(r)} = \mathrm{diag}(\mathbf{B}_r\mathbf{1}_m) \in \mathbb{R}^{n\times n}$ and $\mathbf{D}_m^{(r)} = \mathrm{diag}(\mathbf{B}_r^\top\mathbf{1}_n) \in \mathbb{R}^{m\times m}$. We form the normalized bipartite graph

$$\mathbf{P}_r = (\mathbf{D}_n^{(r)})^{-1/2}\,\mathbf{B}_r\,(\mathbf{D}_m^{(r)})^{-1/2} \in \mathbb{R}^{n\times m},$$

whose entry $[\mathbf{P}_r]_{ij}$ represents the normalized connection strength between sample i and anchor j. Using $\mathbf{P}_r$, we compute the single-view filter

$$\mathbf{G}_r = (1-\alpha)\mathbf{P}_r\big(\mathbf{I}_m - \alpha\mathbf{P}_r^\top\mathbf{P}_r\big)^{-1}\mathbf{P}_r^\top, \qquad (1)$$

where $\alpha \in [0,1)$ controls the diffusion strength. In practice, we choose α so that $\mathbf{I}_m - \alpha\mathbf{P}_r^\top\mathbf{P}_r$ is invertible and numerically stable, e.g., $\alpha < 1/\lambda_{\max}(\mathbf{P}_r^\top\mathbf{P}_r)$.

To clarify the effect of filtering, consider the eigendecomposition $\mathbf{P}_r\mathbf{P}_r^\top = \mathbf{U}\,\mathrm{diag}(\lambda)\,\mathbf{U}^\top$ with $\lambda \geq 0$. The filter applies the scalar transform

$$f(\lambda) = \frac{(1-\alpha)\lambda}{1-\alpha\lambda}, \qquad g(\lambda) = \frac{f(\lambda)}{\lambda} = \frac{1-\alpha}{1-\alpha\lambda}, \qquad (2)$$

where $g(\lambda)$ is increasing on $\lambda \in [0, 1/\alpha)$. Thus, components associated with larger eigenvalues receive a larger relative gain, while small-eigenvalue components are comparatively suppressed. This behavior tends to emphasize dominant graph structure and reduce sensitivity to weak, noisy connectivity.

We combine view-wise filters by a convex mixture

$$\mathbf{A} = \sum_{r=1}^{v}\gamma_r\mathbf{G}_r, \qquad \gamma_r \geq 0,\ \sum_{r=1}^{v}\gamma_r = 1, \qquad (3)$$

which serves as the shared diffusion operator for forming a robust consensus representation.

3.2. Balance-Promoted Graph Fairness Regularization

Filtering reduces representation-level bias, but it does not prevent a protected group from collapsing into a few clusters. To encourage a spread-out allocation inside each group, we add a group-wise alignment term between the continuous embedding and the discrete assignment.

For group g, let $\mathbf{Y}_g \in \{0,1\}^{n_g\times c}$ be the one-hot assignment matrix and let $\mathbf{F}_g = \mathbf{F}_{\mathcal{I}_g,:}$ be the sub-matrix of the global embedding for samples in group g. We use the alignment score $\mathrm{Tr}(\mathbf{F}_g^\top\mathbf{Y}_g)$, which measures how well the embedding rows match their assigned cluster indicators.

This alignment is naturally coupled with the discrete distribution of group g over clusters. Let $n_{g,k}$ denote the number of samples from group g assigned to cluster k. Since $\mathbf{Y}_g$ is one-hot, we have

$$\mathbf{Y}_g^\top\mathbf{Y}_g = \mathrm{diag}(n_{g,1}, \ldots, n_{g,c}), \qquad (4)$$

and therefore the nuclear norm satisfies

$$\|\mathbf{Y}_g\|_* = \sum_{k=1}^{c}\sqrt{n_{g,k}}. \qquad (5)$$

Moreover, because $\mathbf{F}^\top\mathbf{F} = \mathbf{I}_c$, the spectral norm of $\mathbf{F}$ is $\|\mathbf{F}\|_2 = 1$. Writing $\mathbf{F}_g = \mathbf{S}_g\mathbf{F}$ with a row-selection matrix $\mathbf{S}_g$, we have $\|\mathbf{F}_g\|_2 \leq \|\mathbf{F}\|_2 = 1$. By von Neumann's trace inequality,

$$\mathrm{Tr}(\mathbf{F}_g^\top\mathbf{Y}_g) \leq \|\mathbf{F}_g\|_2\,\|\mathbf{Y}_g\|_* \leq \|\mathbf{Y}_g\|_* = \sum_{k=1}^{c}\sqrt{n_{g,k}}. \qquad (6)$$

Since $\sqrt{\cdot}$ is concave, the sum $\sum_{k=1}^{c}\sqrt{n_{g,k}}$ is larger when the counts $\{n_{g,k}\}$ are more even. In addition, by Cauchy–

Schwarz,

$$\sum_{k=1}^{c} \sqrt{n_{g,k}} \leq \sqrt{c \sum_{k=1}^{c} n_{g,k}} = \sqrt{c\, n_g}. \tag{7}$$

We then define the group-normalized regularizer

$$\mathcal{R}_{\text{fair}} = \sum_{g=1}^{G} \frac{\eta}{\sqrt{n_g}} \text{Tr}(\mathbf{F}_g^\top \mathbf{Y}_g), \tag{8}$$

where η controls the trade-off. The factor $1/\sqrt{n_g}$ prevents large groups from dominating; using the bound above, the contribution of each group is upper bounded by $\eta\sqrt{c}$. In optimization, this term discourages a group from concentrating on only a few clusters.

3.3. Unified Optimization

The three parts are integrated into a single joint optimization:

$$\begin{aligned} \max_{\theta} \quad & \text{Tr}\left(\mathbf{F}^\top \mathbf{A} \left(\sum_{r=1}^{v} \beta_r \mathbf{H}_r \mathbf{W}_r\right)\right) + \mathcal{R}_{\text{fair}} \\ \text{s.t.} \quad & \mathbf{F}^\top \mathbf{F} = \mathbf{I}_c, \quad \mathbf{W}_r^\top \mathbf{W}_r = \mathbf{I}_c, \ \forall r, \end{aligned} \tag{9}$$

where $\theta = \{\mathbf{F}, \mathbf{Y}, \{\mathbf{W}_r\}_{r=1}^v, \{\beta_r\}_{r=1}^v, \{\gamma_r\}_{r=1}^v\}$. We enforce the standard feasibility conditions: $\mathbf{Y} \in \{0,1\}^{n\times c}$ with $\mathbf{Y}\mathbf{1}_c = \mathbf{1}_n$, and the fusion/filter weights lie on the probability simplex, i.e., $\beta_r \geq 0$, $\sum_{r=1}^v \beta_r = 1$ and $\gamma_r \geq 0$, $\sum_{r=1}^v \gamma_r = 1$. The diffusion operator $\mathbf{A}$ stabilizes the consensus embedding under noisy views, while $\mathcal{R}_{\text{fair}}$ discourages within-group concentration on a small number of clusters. The discrete fairness metric is enforced in the $\mathbf{Y}$-update step (Sec. 4.4).

4. Optimization

The joint objective in Eq. 9 is non-convex with respect to all variables simultaneously. We therefore adopt an alternating optimization scheme, a common strategy for constrained multi-variable problems, in which one block of variables is updated at a time while keeping the others fixed. Each subproblem admits a closed-form or efficiently computable solution, ensuring convergence to a high-quality local optimum in practice. The overall procedure is summarized in Algorithm 1.

4.1. Update of Alignment Matrices $\{\mathbf{W}_r\}$

Fixing $\mathbf{F}$, $\mathbf{Y}$, β_r and γ_r, the optimization for $\mathbf{W}_r$ becomes:

$$\max_{\mathbf{W}_r} \text{Tr}\left(\mathbf{W}_r^\top \left[\beta_r \mathbf{H}_r^\top \mathbf{A}^\top \mathbf{F}\right]\right), \quad \text{s.t. } \mathbf{W}_r^\top \mathbf{W}_r = \mathbf{I}_c, \tag{10}$$

where $\mathbf{A} = \sum_{j=1}^v \gamma_j \mathbf{G}_j$. This is a classic Orthogonal Procrustes problem, whose closed-form solution is obtained by computing the SVD $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ of the coefficient matrix and setting $\mathbf{W}_r = \mathbf{U}\mathbf{V}^\top$.

Algorithm 1 FLFMVC Optimization Procedure

Input: Base partitions $\{\mathbf{H}_r\}_{r=1}^v$, number of clusters c, group index sets $\{\mathcal{I}_g\}_{g=1}^G$, hyperparameters η, α, m, τ (and k for k-NN if used). **Output**: Final consensus assignment $\mathbf{Y}$.

1: Construct fairness-aware graph filters $\{\mathbf{G}_r\}_{r=1}^v$ via proportional anchor sampling.
2: Initialize $\mathbf{Y}^{(0)}$ (e.g., by k-means on row-normalized $[\mathbf{H}_1, \ldots, \mathbf{H}_v]$), $\mathbf{F}$, $\{\mathbf{W}_r\}$, and simplex weights $\boldsymbol{\beta}, \boldsymbol{\gamma}$.
3: **repeat**
4: **Update Alignment:** Solve for each alignment matrix $\mathbf{W}_r$ via Orthogonal Procrustes (Sec. 4.1).
5: **Update Weights:** Update view and filter weights $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ adaptively based on their contribution to the objective using a softmax scheme (Sec. 4.2).
6: **Update Embedding:** Update the consensus embedding $\mathbf{F}$ by solving the Orthogonal Procrustes problem incorporating both the robustness and fairness terms (Sec. 4.3).
7: **Update Discretization:** Refine the discrete assignment $\mathbf{Y}$ from $\mathbf{F}$ using the fairness-guided strategy (Sec. 4.4).
8: **until** convergence criteria are met
9: **return** $\mathbf{Y}$.

Numerical stability. The SVDs in our updates are thin decompositions of at most $n \times c$ (or $c \times c$) matrices with small c, and the Procrustes solution $\mathbf{U}\mathbf{V}^\top$ remains well-defined even under rank deficiency (zero singular values do not affect $\mathbf{U}\mathbf{V}^\top$). When the coefficient matrix is nearly singular, a standard remedy is to add a tiny ridge term (e.g., $\epsilon\mathbf{I}$) to improve conditioning, which we found sufficient in practice.

4.2. Update of Weights $\{\beta_r\}$ and $\{\gamma_r\}$

The view-fusion weights $\{\beta_r\}$ and filter-combination weights $\{\gamma_r\}$ are updated to reflect each view's contribution in the current iteration. To ensure the weights are non-negative and sum to one, we employ a softmax function.

For β_r, the contribution score is first computed as:

$$c_{\beta_r} = \text{Tr}\left(\mathbf{F}^\top \mathbf{A} \mathbf{H}_r \mathbf{W}_r\right). \tag{11}$$

The weight vector $\boldsymbol{\beta} = [\beta_1, \ldots, \beta_v]^\top$ is then updated via the softmax function:

$$\beta_r \leftarrow \frac{\exp(\kappa \cdot c_{\beta_r})}{\sum_{j=1}^v \exp(\kappa \cdot c_{\beta_j})}, \tag{12}$$

where $\kappa > 0$ is a temperature parameter that controls the sharpness of the distribution. A similar procedure is applied to update γ_r using the scores $c_{\gamma_r} = \text{Tr}(\mathbf{F}^\top \mathbf{G}_r \mathbf{B})$ with $\mathbf{B} = \sum_{j=1}^v \beta_j \mathbf{H}_j \mathbf{W}_j$. This softmax normalization

not only guarantees non-negativity but also adaptively allocates influence among views, preventing any single view from dominating the optimization. In our experiments, we set $\kappa = 1$ and found the results to be robust.

4.3. Update of Consensus Embedding $\mathbf{F}$

With other variables fixed, the update for $\mathbf{F}$ maximizes:

$$\mathrm{Tr}\big(\mathbf{F}^\top \mathbf{M}_\mathbf{A}\big) + \sum_{g=1}^{G} \frac{\eta}{\sqrt{n_g}} \mathrm{Tr}(\mathbf{F}_g^\top \mathbf{Y}_g), \tag{13}$$

where $\mathbf{M}_\mathbf{A} = (\sum_{r=1}^{v} \gamma_r \mathbf{G}_r)(\sum_{r=1}^{v} \beta_r \mathbf{H}_r \mathbf{W}_r)$ encodes the robustness component, and the second term encodes the fairness regularizer. Let $\mathbf{M}_\mathbf{B}$ be the block matrix constructed from $\{\mathbf{Y}_g\}$ and group weights $\frac{\eta}{\sqrt{n_g}}$. The subproblem becomes:

$$\max_{\mathbf{F}^\top \mathbf{F} = \mathbf{I}_c} \mathrm{Tr}\big(\mathbf{F}^\top (\mathbf{M}_\mathbf{A} + \mathbf{M}_\mathbf{B})\big), \tag{14}$$

which is again an Orthogonal Procrustes problem solved via SVD of $(\mathbf{M}_\mathbf{A} + \mathbf{M}_\mathbf{B})$.

4.4. Update of Discrete Assignments $\mathbf{Y}$

The discrete assignments $\mathbf{Y}$ are obtained from $\mathbf{F}$ through our fairness-guided discretization strategy. Let Δ_i denote the margin between the largest and second-largest entries in the i-th row of $\mathbf{F}$. If $\Delta_i \geq \tau$, we assign the sample to the argmax cluster; otherwise, we evaluate all feasible assignments and select the one that maximally improves the discrete fairness metric (e.g., f_{CCE}) while preserving balance constraints. This ensures that fairness optimization is aligned with the final discrete solution, rather than only its continuous relaxation.

4.5. Computational Complexity

Per iteration, FLFMVC costs $\mathcal{O}\big(v\, n\, c^2 + v\, c^3\big)$, which is near-linear in n for fixed (v, c): the n-dependence comes from forming/filtering products and the $n \times c$ Procrustes SVD (both $\mathcal{O}(nc^2)$), while the v view-wise Orthogonal Procrustes updates contribute $\mathcal{O}(vc^3)$. Using the bipartite factorization, applying each graph filter is $\mathcal{O}(vnm)$ and, under the common setting $m = \Theta(c)$, is absorbed into the $\mathcal{O}(vnc^2)$ term. The only superlinear (in n) operation is independent of n: computing $(\mathbf{I}_m - \alpha \mathbf{P}_r^\top \mathbf{P}_r)^{-1}$ per view is a *one-off* initialization cost of $\mathcal{O}(v\, m^3)$, which does not appear in the per-iteration bound. Therefore, the overall runtime is dominated by $\mathcal{O}(vnc^2 + vc^3)$ per iteration, plus a single $\mathcal{O}(vm^3)$ setup.

5. Experiments

We conduct extensive experiments to evaluate our proposed FLFMVC framework. Our evaluation is designed to answer several key questions: its performance against state-of-the-art methods, the contribution of its core components, and its utility on a real-world application.

5.1. Experimental Setup

5.1.1 Datasets.

We evaluate our method on 9 diverse benchmark datasets, summarized in Table 1. This collection includes native multi-view datasets (HAR, PAMAP2) and widely-used single-view datasets for which we generate multiple views using a standard multi-kernel strategy. The datasets span various domains, including human activity recognition, medical diagnosis, and census data, and feature a wide range of sample sizes, cluster counts, and protected attributes.

Table 1. Summary of the benchmark datasets.

Dataset	Sample	Feature	Cluster	Protected Group
Adult	32,561	107	2	Race (5)
HAR	10,299	9 views	6	Person ID (30)
Heart Disease	303	26	5	Sex (2)
Jaffe	213	676	10	Expression (7)
Obesity	2,111	31	7	Sex (2)
PAMAP2	18,716	7 views	7	Age (3)
Student-Mat	395	59	2	Sex (2)
Student-Por	649	59	2	Sex (2)
Yale	165	1,024	15	Glasses (2)

5.1.2 Baselines.

We compare FLFMVC against ten representative algorithms from the literature discussed in Section 2. The baselines are selected to cover the state-of-the-art in late fusion, fairness, and balance. For a fair comparison, all methods are tuned following their authors' recommended procedures. The selected baselines are:

- **Late Fusion MVC Baselines**: OPLF [9], ALMVC [14], M3LF [7], LFLKA [10], sLGm [12], and CFGFLF [3].
- **Fair and Balanced Clustering Baselines**: SpFC [6], FMSC [8], KFC [4], and FCE [16].

5.1.3 Evaluation Metrics.

We evaluate clustering quality using Accuracy (ACC) and Normalized Mutual Information (NMI). To assess the trade-off between fairness and balance, we adopt the rigorous **f_CCE** metric from Zhou et al. [16]. This metric is designed as a stringent worst-case measure. Given a clustering result $C = \{\pi_1, \ldots, \pi_c\}$ and a set of protected groups

Table 2. Performance comparison with late fusion baselines across ACC and NMI.

Method	Adult	HAR	Heart	Jaffe	Obesity	PAMAP2	Student-Mat	Student-Por	Yale
	ACC								
LFLKA [TMM, 2021]	0.7177	0.5955	0.3568	0.9516	0.4353	0.5158	0.5511	0.6672	0.5127
OPLF [ICML, 2021]	0.7152	0.3475	0.3462	0.9681	0.4323	0.1730	0.5089	0.6271	0.5170
sLGm [SP, 2022]	0.7164	0.4641	0.3439	0.9380	0.4358	0.5107	0.5491	0.6598	0.5139
ALMVC [ACM MM, 2022]	0.7151	**0.6052**	0.3521	0.9272	0.4353	0.4814	0.5476	0.6635	0.4806
M3LF [TNNLS, 2024]	0.7159	0.5262	0.3350	0.9089	0.3963	0.6485	0.5646	0.6687	0.4733
CFGFLF [IEEE SPL, 2025]	0.7504	0.4582	0.5149	0.8230	0.4301	0.2765	**0.6709**	0.8459	**0.5533**
FLFMVC (Ours)	**0.7589**	0.5270	**0.5178**	**0.9948**	**0.4368**	**0.7735**	**0.6709**	**0.8490**	0.5303
	NMI								
LFLKA [TMM, 2021]	**0.1239**	0.4712	0.1263	0.9408	0.3007	0.3725	**0.0069**	0.0698	0.5315
OPLF [ICML, 2021]	0.1197	0.2160	**0.1439**	0.9484	0.3000	0.0003	0.0002	0.0516	0.5357
sLGm [SP, 2022]	0.1196	0.3580	0.1378	0.9398	0.2993	0.4044	0.0057	0.0640	0.5324
ALMVC [ACM MM, 2022]	0.1189	**0.4809**	0.1261	0.9268	0.2991	0.3514	0.0054	**0.0742**	0.5225
M3LF [TNNLS, 2024]	0.1188	0.3726	0.1314	0.9141	0.2654	0.6379	0.0065	0.0625	0.5006
CFGFLF [IEEE SPL, 2025]	0.0005	0.3278	0.0129	0.8977	0.2934	0.1123	0.0044	0.0066	0.5532
FLFMVC (Ours)	0.0109	0.3991	0.0329	**0.9908**	**0.3023**	**0.7183**	0.0010	0.0135	**0.5558**

Table 3. Performance comparison with fairness baselines across ACC, NMI and f_CCE.

Method	Adult	HAR	Heart	Jaffe	Obesity	PAMAP2	Student-Mat	Student-Por	Yale
	ACC								
SpFC [ICML, 2019]	**0.7595**	0.2129	**0.4785**	0.9437	0.3979	0.1732	**0.6405**	**0.7350**	0.2000
KFC [NeurIPS, 2020]	0.5696	0.4024	0.3927	0.3474	0.2293	0.3319	**0.6405**	0.5054	0.3030
FMSC [ACM MM, 2024]	0.6579	0.2492	0.3894	0.1408	0.2122	0.2288	0.5316	0.5901	0.2364
FCE [TPAMI, 2025]	0.6173	**0.5397**	0.3069	0.6761	0.3515	0.4315	**0.6405**	**0.7350**	0.4000
FLFMVC (Ours)	0.6700	0.5270	0.3155	**0.9948**	**0.4026**	**0.7428**	0.5828	0.6040	**0.5303**
	NMI								
SpFC [ICML, 2019]	0.0009	0.0350	**0.1549**	0.9334	**0.2608**	0.0004	0.0050	0.0565	0.1933
KFC [NeurIPS, 2020]	0.0570	0.2306	0.1158	0.3567	0.0603	0.2753	0.0001	0.0188	0.3571
FMSC [ACM MM, 2024]	0.0072	0.1129	0.0539	0.0122	0.1085	0.0764	0.0017	0.0045	0.2546
FCE [TPAMI, 2025]	0.0579	0.3552	0.0797	0.7084	0.2140	0.3418	0.0022	0.0473	0.4656
FLFMVC (Ours)	**0.0938**	**0.3991**	0.1179	**0.9908**	0.2454	**0.6840**	**0.0229**	**0.0655**	**0.5558**
	f_CCE								
SpFC [ICML, 2019]	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.5714	0.0000
KFC [NeurIPS, 2020]	0.7623	0.0000	0.0971	0.2583	0.0201	0.0052	0.1497	0.8460	0.0000
FMSC [ACM MM, 2024]	0.0000	0.1210	0.0000	0.0000	0.0000	0.0000	0.9412	0.0000	0.0000
FCE [TPAMI, 2025]	0.8266	0.0000	0.2062	0.0000	0.1007	0.2653	**0.9947**	0.8460	0.0000
FLFMVC (Ours)	**0.9260**	**0.1581**	**0.4358**	**0.6452**	**0.3548**	**0.4736**	**0.9947**	**0.9974**	**0.1795**

$G = \{G_1, \ldots, G_G\}$, the metric is built upon two core components:

Accuracy (ACC). Accuracy measures the percentage of correctly clustered samples by finding the best one-to-one mapping between the predicted cluster labels and the ground-truth class labels. Given a set of n samples, let l_i be the ground-truth label and c_i be the predicted cluster label for the i-th sample. ACC is defined as:

$$\mathrm{ACC} = \frac{1}{n} \sum_{i=1}^{n} \delta(\mathrm{map}(c_i), l_i), \tag{15}$$

where $\delta(x, y)$ is the indicator function, which is 1 if $x = y$ and 0 otherwise. The function $\mathrm{map}(\cdot)$ represents the optimal permutation mapping between cluster labels and ground-truth labels, which is found efficiently using the

Hungarian algorithm.

Normalized Mutual Information (NMI). Normalized Mutual Information evaluates the quality of clustering by measuring the mutual dependence between the predicted clustering and the ground-truth classes, adjusted for chance. Let $\Omega = \{\omega_1, \ldots, \omega_K\}$ be the set of ground-truth classes and $C = \{c_1, \ldots, c_J\}$ be the set of predicted clusters. NMI is defined as:

$$\mathrm{NMI}(\Omega, C) = \frac{2 \times I(\Omega; C)}{H(\Omega) + H(C)}, \tag{16}$$

where $I(\Omega; C)$ is the mutual information between the two partitions, and $H(\Omega)$ and $H(C)$ are the entropies of the ground-truth and predicted partitions, respectively. The mutual information and entropy are calculated as:

$$I(\Omega; C) = \sum_{k=1}^{K} \sum_{j=1}^{J} \frac{|\omega_k \cap c_j|}{n} \log \frac{n\,|\omega_k \cap c_j|}{|\omega_k|\,|c_j|}, \tag{17}$$

$$H(\Omega) = -\sum_{k=1}^{K} \frac{|\omega_k|}{n} \log \frac{|\omega_k|}{n}. \tag{18}$$

An NMI score of 1 indicates a perfect match, while a score of 0 indicates that the two partitions are completely independent.

Fairness_Cluster-Capacity-Equality (f_CCE). First, the **Cluster-level Fairness Score** (F_k) for a single cluster π_k evaluates its fairness across all groups. It is defined as the worst-case fairness score among all protected groups:

$$F_k = \min_{g \in G} \left(\min \left(c\,\gamma_{g,k}, \frac{1}{c\,\gamma_{g,k}} \right) \right), \tag{19}$$

where $\gamma_{g,k} = |\pi_k \cap G_g| / |G_g|$ is the proportion of group g assigned to cluster π_k.

Second, the **Global Cluster Balance Score** (B_C) measures the size uniformity of the entire clustering:

$$B_C = \frac{\min_{i \in \{1,\ldots,c\}} |\pi_i|}{\max_{j \in \{1,\ldots,c\}} |\pi_j|}. \tag{20}$$

The final f_CCE score is then defined as the minimum harmonic mean of these two components over all clusters:

$$\text{f\_CCE}(C) = \min_{k \in \{1,\ldots,c\}} \left(2 \cdot \frac{F_k \cdot B_C}{F_k + B_C} \right). \tag{21}$$

This formulation rigorously ensures that a high score is achieved only if *every* cluster is fair to *all* groups, and the overall clustering is balanced. For all metrics, higher values indicate better performance.

5.1.4 Implementation Details

All experiments were conducted in MATLAB R2022b on a machine with an Intel Core i7-7700 CPU and 64GB of RAM. For single-view datasets, we generated multiple views using a standard multi-kernel strategy with adaptively determined Gaussian kernels.

In our implementation, $\mathbf{H}_r$ is obtained via Spectral Kernel k-Means (SKKM). For single-view datasets, we first generate v kernel matrices $\{\mathbf{K}_r\}_{r=1}^{v}$ with different kernel types/parameters, and then run SKKM on each $\mathbf{K}_r$ to obtain $\mathbf{H}_r \in \mathbb{R}^{n \times c}$ (the top-c eigenvectors of the symmetrized kernel matrix). For native multi-view datasets, SKKM is applied view-wisely to each view to produce $\mathbf{H}_r$ with the same target dimension c. Hyperparameters for all methods were tuned via grid search; detailed search space and hyperparameter analysis are provided in the supplementary material. For our FLFMVC, the ambiguity threshold τ was fixed at 0.1 for all main experiments, a choice justified by subsequent sensitivity analysis. All reported results are the average of 10 independent runs with different random seeds. Our code is provided in the supplementary material.

5.2. Comparison with State-of-the-Art Methods

We now present the main experimental results. Given the distinct design philosophies of the baseline, some prioritizing accuracy, others fairness, we structure our comparison into two parts for a more insightful evaluation. For all comparison methods, we follow their recommended parameter configurations and search methodologies to ensure a fair comparison. The best results among this group are in **bold** and the second-best are underlined.

5.3. Hyperparameter Settings and Analysis

This section details the hyperparameter configurations used for our FLFMVC in the main experiments and provides a summarized analysis of our findings from the grid search process. Our goal is to offer insights into the model's behavior and provide guidance for future applications.

5.3.1 Hyperparameter Configurations

As described in the main paper, our experiments are structured into two comparison tracks with different objectives. Consequently, the optimal hyperparameters for FLFMVC also differ. Table 4 lists the parameters used for comparison against late fusion baselines (optimizing for ACC). Table 5 lists the parameters for comparison against fairness baselines (optimizing for f_CCE).

5.3.2 Analysis of Hyperparameter Choices

Our extensive grid search revealed distinct patterns in the optimal hyperparameter choices across different data

Table 4. FLFMVC hyperparameter settings for the "performance track" (see 2), selected to maximize ACC.

Dataset	Emb. Factor	α	m	η	k-NN
Adult	4	0.99	8	1e-5	5
HAR	5	0.1	100	1e+1	5
Heart	4	0.99	20	1e-3	5
Jaffe	5	0.95	80	1e+0	5
Obesity	5	0.1	14	1e-1	5
PAMAP2	5	0.99	56	1e-4	5
Student-Mat	2	0.7	4	1e-5	5
Student-Por	3	0.7	4	1e-5	5
Yale	1	0.95	120	1e+0	5

Table 5. FLFMVC hyperparameter settings for the "fairness track" (see 3), selected to maximize f_CCE.

Dataset	Emb. Factor	α	m	η	k-NN
Adult	1	0.1	8	1e-5	5
HAR	5	0.1	100	1e+1	5
Heart	5	0.95	40	1e+0	10
Jaffe	5	0.95	80	1e+0	5
Obesity	1	0.99	25	1e+0	5
PAMAP2	5	0.99	14	1e-5	5
Student-Mat	3	0.95	8	1e-1	5
Student-Por	3	0.99	16	1e-5	5
Yale	1	0.95	120	1e+0	5

modalities, which we summarize below.

Image Datasets (Jaffe, Yale). These datasets consistently favored a high filter parameter α ($\geq$ 0.95), suggesting that emphasizing the principal graph structure is beneficial for clustering image data. The optimal fairness trade-off η, however, varied, indicating that the complexity of the fairness objective is highly data-specific. For new image datasets, we recommend prioritizing a high α while exploring a wide, logarithmic range for η.

Tabular Datasets (Adult, Heart, etc.). When optimizing for ACC, this diverse group of datasets often preferred a small η (e.g., $\leq 10^{-3}$), focusing the optimization almost entirely on the clustering objective. The optimal α, however, showed no clear trend, being highly data-dependent. Notably, achieving a high f_CCE score, as seen on the Obesity dataset, sometimes required a completely different parameter regime (e.g., high η and high α) compared to optimizing for ACC. This "dual characteristic" highlights the importance of a thorough grid search on new tabular data.

Multi-view Sensor Datasets (HAR, PAMAP2). This category exhibited the most unique and consistent profile, strongly favoring a high fairness trade-off η (e.g., $\geq$ 10) and a low filter parameter α (e.g., $\leq$ 0.1). We hypothesize that a high η is necessary to effectively navigate the complex, fused feature space, while a low α helps maintain inter-view consistency. For new multi-view sensor datasets, we strongly recommend starting the search in this specific "high-η, low-α" region.

Performance Against Late Fusion Baselines. Table 2 places FLFMVC in a head-to-head comparison with state-of-the-art late fusion methods, where the primary objective is maximizing clustering quality.

The results show that FLFMVC is highly competitive even on this traditional battleground. It achieves the best or second-best ACC on 8 out of nine datasets, notably securing the top performance on complex, high-dimensional datasets like Jaffe and PAMAP2. This demonstrates that our integrated fairness mechanisms do not compromise the core clustering capability of the robust late fusion framework. In several cases, such as on Heart and PAMAP2, our method substantially outperforms the strong CFGFLF baseline, suggesting that our fairness-aware filter can also lead to a better consensus representation.

Performance Against Fairness Baselines. Table 3 evaluates FLFMVC against specialized fair clustering methods. We report the results for each method at the hyperparameter setting that maximizes its f_CCE score, provided the corresponding ACC is not trivially low. The results are unequivocal: our FLFMVC is in a class of its own, achieving the best f_CCE score on all nine datasets, often by a remarkable margin. This dominance confirms the effectiveness of our multi-level fairness mechanism.

Furthermore, FLFMVC achieves this state-of-the-art fairness while maintaining high clustering accuracy. Some baselines, like SpFC on Adult and Heart, can achieve a high ACC, but this comes at the complete expense of fairness, resulting in an f_CCE score of 0.0000. They are forced into a harsh trade-off. Our FLFMVC, in contrast, consistently finds a superior balance. For instance, on Jaffe, it simultaneously achieves the best ACC, NMI, and f_CCE, showcasing its ability to resolve the accuracy-fairness dilemma that constrains prior work.

Interpreting the Performance Trade-offs. A deeper analysis of Tables 2 and 3 reveals a fundamental pattern distinguishing FLFMVC from prior work. The fairness-centric baselines, such as SpFC and FCE, often exhibit a stark accuracy-fairness trade-off: achieving a high f_{CCE} score frequently comes at the cost of a significant drop in ACC, as their fairness constraints can rigidly conflict with the data's inherent cluster structure. Conversely, the performance-focused LFMVC methods achieve high accuracy by remaining completely agnostic to fairness, often resulting in near-zero f_{CCE} scores on socially sensitive datasets.

FLFMVC, in contrast, consistently operates in a more desirable "win-win" region of the performance space. We attribute this to our multi-level design. The fairness-aware filter creates a more robust and less biased representation to begin with, reducing the burden on the subsequent optimization. The group-normalized regularizer then gently guides the solution towards an equitable state, rather than imposing a hard constraint. This synergy allows our framework to discover cluster structures that are not only accurate but also intrinsically fair and balanced.

5.4. Ablation Studies

To dissect the contribution of our framework's key components, we conduct a detailed ablation study on the challenging, large-scale PAMAP2 dataset. For a fair comparison, all variants are run using the same fixed hyperparameters as our main model for this dataset. The results are presented in two focused analyses.

Table 6. Ablation on the effect of graph filtering on PAMAP2. N/A indicates that the method was not run due to prohibitive computational cost.

Metric	$\mathbf{M}_{base}$	$\mathbf{M}_{base}$+$\mathbf{F}_{full}$	$\mathbf{M}_{base}$+$\mathbf{F}_{fair}$	FLFMVC (Full)
ACC	0.5333	N/A	**0.7540**	0.7428
NMI	0.4055	N/A	**0.7025**	0.6840
f_CCE	0.1414	N/A	0.0178	**0.4736**

Table 7. Ablation on the effect of the fairness components on PAMAP2.

Metric	$\mathbf{M}_{base}$	$\mathbf{M}_{base}$+$\mathbf{R}_A$	$\mathbf{M}_{base}$+$\mathbf{R}_F$	FLFMVC (Full)
ACC	0.5333	0.5333	0.5332	**0.7428**
NMI	0.4055	0.4055	0.4052	**0.6840**
f_CCE	0.1414	0.1414	0.2011	**0.4736**

Efficacy of Graph Filtering. Table 6 investigates the impact of our graph filtering mechanism. The results show that filtering is crucial for handling this complex, multi-view dataset. Our baseline model ($\mathbf{M}_{base}$), which relies solely on the late-fusion alignment, achieves a modest ACC of 0.5333. By incorporating our proposed fairness-aware filter ($\mathbf{M}_{base}$+$\mathbf{F}_{fair}$), the ACC score dramatically improves to 0.7540, a gain of over 22 percentage points. This underscores the filter's powerful ability to learn a robust consensus representation from noisy base partitions. The standard full-graph filter ($\mathbf{M}_{base}$+$\mathbf{F}_{full}$) was omitted from the comparison as its runtime was prohibitively long on this large-scale dataset. This validates the practical necessity and efficiency of our bipartite graph approach. Our full model (**FLFMVC**) retains this strong accuracy while simultaneously enhancing fairness, resulting in a slight, expected trade-off in ACC compared to the variant focused solely on filtering.

Efficacy of Fairness Components. Table 7 isolates the contribution of our objective-level fairness mechanisms. We compare four variants: the baseline without any fairness components ($\mathbf{M}_{base}$); the baseline with our fairness regularizer but using a standard argmax discretization ($\mathbf{M}_{base}$+$\mathbf{R}_A$); the baseline with the regularizer and our proposed fairness-guided discretization ($\mathbf{M}_{base}$+$\mathbf{R}_F$); and our complete model (**FLFMVC (Full)**).

The results reveal a critical finding: simply adding the regularizer ($\mathbf{M}_{base}$+$\mathbf{R}_A$) yields no improvement in the f_CCE score, as the naive discretization step nullifies its benefits. In contrast, activating our fairness-guided discretization ($\mathbf{M}_{base}$+$\mathbf{R}_F$) significantly jumps the f_CCE score from 0.1414 to 0.2011, demonstrating its effectiveness. The final synergy is shown by our full model, **FLFMVC**, which integrates both the fair filter and the complete fairness module. This combination achieves the best performance across all metrics, boosting f_CCE to a state-of-the-art 0.4736 while also attaining the highest accuracy. This confirms that our innovations at both the representation and objective levels are critical and work synergistically.

5.5. Accuracy-Fairness Trade-off Analysis

To visualize the trade-off between clustering quality and fairness, Figure 1 plots the Pareto fronts (ACC vs. f_CCE) for our method and key baselines, with the top-right corner representing ideal performance. A clear pattern emerges across all datasets: the Pareto front for our FLFMVC consistently dominates the alternatives. This demonstrates a relationship of Pareto dominance, as FLFMVC offers a superior accuracy-fairness profile. For instance, compared to a specialized method like FCE, which often pays a heavy price in accuracy for its fairness gains, FLFMVC consistently maintains a high accuracy across a wider range of fairness values. This visual evidence confirms that our framework offers a fundamentally better compromise, pushing the boundary of what is achievable in fair multi-view clustering.

5.6. Convergence Analysis

We conclude our analysis by examining the convergence behavior of our optimization strategy, illustrated in Figure 2 for two representative datasets. The objective function value exhibits a sharp increase within the first few iterations before rapidly stabilizing. Our algorithm consistently meets its convergence criterion, a relative objective increase between iterations of less than 10^{-4}, in fewer than 10 iterations on both datasets. This rapid convergence demonstrates the stability and efficiency of our optimization approach.

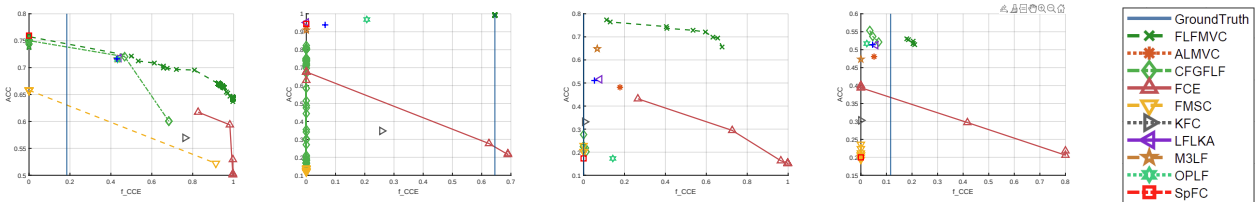

(a) Adult (b) Jaffe (c) PAMAP2 (d) Yale

Figure 1. Accuracy (ACC) vs. Fairness (f_CCE) Trade-off: Pareto Front Analysis on Adult, Jaffe, PAMAP2, and Yale datasets. The shared legend is shown on the right. The vertical blue line represents the fairness score of the ground-truth partition.

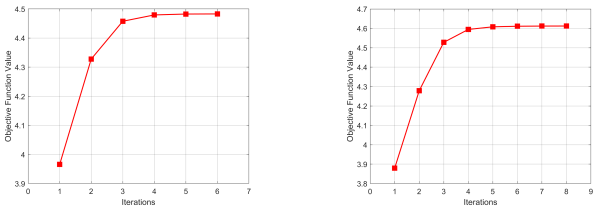

(a) Adult (b) Student-Mat

Figure 2. Convergence of the FLFMVC algorithm on the (a) Adult and (b) Student-Mat datasets.

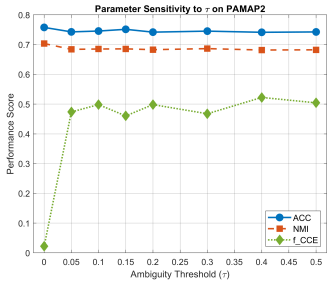

Figure 3. Sensitivity analysis of the ambiguity threshold τ on the PAMAP2 dataset.

5.7. Parameter Sensitivity Analysis

To address the concern of our framework's reliance on the ambiguity threshold τ, we analyze its sensitivity on the PAMAP2 dataset (Figure 3). The results demonstrate a remarkable robustness. Setting $\tau = 0$ (a standard argmax) yields a poor f_CCE score, but the fairness metric immediately jumps and then remains high and stable across a broad [0.05, 0.5] range. Critically, both ACC and NMI are almost entirely unaffected by the choice of τ in this range. This confirms our method's insensitivity to this parameter and validates our choice of $\tau = 0.1$ for the main experiments.

5.8. Application: Fair Person Identification

We demonstrate the practical impact of FLFMVC on the classic JAFFE visual dataset. The task is to cluster facial images into 10 groups, corresponding to the 10 individuals in the dataset. In this scenario, we treat the 7 facial expressions (e.g., 'Happy', 'Angry') as protected attributes. A fair and robust system must identify an individual correctly, irrespective of their emotional expression—a critical requirement for applications ranging from biometric security to empathetic human-computer interaction.

We compare FLFMVC with the fairness-agnostic but high-performing CFGFLF. The results in Figure 4 reveal a clear contrast: the baseline CFGFLF achieves a decent ACC of 0.8977 but fails completely on fairness ($f_{\text{CCE}} = 0.0000$), while our FLFMVC attains both high fairness ($f_{\text{CCE}} = 0.4305$) and near-perfect accuracy (0.9908).

The visualization in Figure 4 explains this outcome. CFGFLF's clusters exhibit severe representational skew. For instance, the clusters for 'Person 2' and 'Person 8' are missing multiple expression categories entirely, indicating the model has failed to learn a generalizable identity for these individuals. It has effectively learned to recognize them only under a limited set of expressions. In contrast, FLFMVC produces clusters where each person's full range of expressions is represented far more equitably. The stacked bars are visually more consistent across all 10 individuals, confirming that our method successfully disentangles personal identity from transient emotional states.

The operational implications for visual systems are significant. By ensuring high fairness, FLFMVC delivers a more reliable and trustworthy identification system. This is achieved through our integrated mechanism: the fairness-aware filter learns a robust representation from the visual features, while the balance regularizer and guided discretization ensure that this representation is not biased towards over-represented expressions. Because FLFMVC operates on base partitions, it can serve as a drop-in, privacy-friendly consensus layer over features extracted by different neural networks. This case study powerfully illustrates that for visual media tasks, enforcing fairness is not merely an ethical constraint but can be a pathway to achieving superior robustness and accuracy.

6. Conclusion

We presented FLFMVC, a late-fusion multi-view clustering framework that combines robustness to noisy views with group fairness and cluster balance. The method integrates a fairness-aware bipartite graph filter, a group-normalized balance regularizer, and a fairness-guided dis-

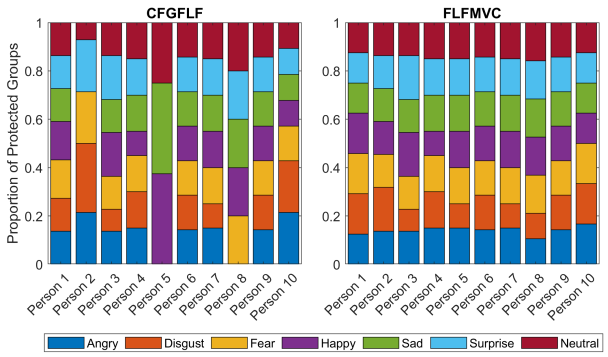

Figure 4. Expression distribution within each person-cluster on the JAFFE dataset. Each bar represents a cluster corresponding to an individual.

cretization step. Across nine benchmarks, FLFMVC achieves higher fairness and better balance while maintaining competitive clustering accuracy.

Future work includes more automatic hyperparameter selection and extensions to larger visual settings such as video and 3D data, especially when protected-attribute annotations are noisy or incomplete.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (62376146,62176001,62472294), the Shanxi Provincial Central Guiding Fund for Local Science and Technology Development Special Project (YDZJSX20231D003).

References

[1] Y. Chen, Z. Friggstad, B. Li, G. L. G. Lortie, A. Ramesh, S. R. Safavi, M. R. Salavatipour, and J. Solomon. Proportionally fair clustering. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning (ICML)*, volume 97 of *Proceedings of Machine Learning Research*, pages 1032–1041, Long Beach, CA, USA, 2019. PMLR. 1

[2] F. Chierichetti, R. Kumar, S. Lattanzi, and S. Vassilvitskii. Fair clustering through fairlets. In *Advances in Neural Information Processing Systems 30 (NeurIPS 2017)*, pages 5029–5037, Red Hook, NY, USA, 2017. Curran Associates, Inc. 1

[3] Y. Guo, H. Jiang, Y. Chen, and L. Du. A scalable consensus fast graph filtering approach for late fusion multi-view clustering. *Signal Processing*, 237:110074, 2025. 2, 3, 5

[4] E. Harb and H. S. Lam. Kfc: A scalable approximation algorithm for k-center fair clustering. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 14509–14519, Red Hook, NY, USA, 2020. Curran Associates, Inc. 1, 2, 5

[5] M. Kleindessner, S. Kumar, P. Awasthi, and J. Morgenstern. Guarantees for spectral clustering with fairness constraints. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning (ICML)*, volume 97 of *Proceedings of Machine Learning Research*, pages 3458–3468, Long Beach, CA, USA, 2019. PMLR. 1

[6] M. Kleindessner, S. Samadi, P. Awasthi, and J. Morgenstern. Guarantees for spectral clustering with fairness constraints. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning (ICML)*, volume 97 of *Proceedings of Machine Learning Research*, pages 3458–3467, Long Beach, CA, USA, 2019. PMLR. 2, 5

[7] M. Li, X. Liu, Y. Zhang, and W. Liang. Late fusion multiview clustering via min-max optimization. *IEEE Transactions on Neural Networks and Learning Systems*, 35(7):9417–9427, 2024. 2, 5

[8] R. Li, H. Hu, L. Du, J. Chen, B. Jiang, and P. Zhou. One-stage fair multi-view spectral clustering. In *Proceedings of the 32nd ACM International Conference on Multimedia (MM '24)*, pages 1–10, Melbourne, VIC, Australia, 2024. ACM. 2, 5

[9] X. Liu, L. Liu, Q. Liao, S. Wang, Y. Zhang, W. Tu, C. Tang, J. Liu, and E. Zhu. One pass late fusion multi-view clustering. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, volume 139 of *Proceedings of Machine Learning Research*, pages 6850–6859, Virtual Event, 2021. PMLR. 1, 2, 5

[10] Y. Liu, T. Zhang, and D. Zhao. Late fusion multiple kernel clustering with local kernel alignment. *IEEE Transactions on Multimedia*, 23:4321–4334, 2021. 2, 5

[11] S. Wang, X. Liu, E. Zhu, C. Tang, J. Liu, J. Hu, J. Xia, and J. Yin. Multi-view clustering via late fusion alignment maximization. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 3778–3784, Macao, China, 2019. IJCAI.org. 2

[12] X. Wu, C. Huang, X. Liu, F. Zhou, and Z. Ren. Multiple kernel clustering with shifted laplacian on grassmann manifold. In *Proceedings of the 32nd ACM International Conference on Multimedia (MM '24)*, pages 2448–2456, Melbourne, VIC, Australia, 2024. ACM. 2, 5

[13] Y. Yang and H. Wang. Multi-view clustering: A survey. *Big Data Mining and Analytics*, 1(2):83–107, 2018. 1

[14] G. Zhao, L. Chen, and H. Wang. Efficient anchor learning-based multi-view clustering: A late fusion method. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 1545–1553, Lisbon, Portugal, 2022. ACM. 2, 5

[15] H. Zhao, Z. Ding, Y. Guo, and Y. Fu. Multi-view clustering via deep matrix factorization with late fusion alignment. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI 2019)*, pages 4102–4108, Macao, China, 2019. International Joint Conferences on Artificial Intelligence Organization. 1

[16] P. Zhou, R. Li, Z. Ling, L. Du, and X. Liu. Fair clustering ensemble with equal cluster capability. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(3):1729–1746, 2025. 2, 5