A General Anchor-Based Framework for Scalable Fair Clustering

Shengfei Wei*[1], Suyuan Liu*[1], Jun Wang[1], Ke Liang[1], Miaomiao Li†[2], Lei Luo†[1]

[1] College of Computer Science and Technology, National University of Defense Technology, Changsha, 410073, China
[2] the College of Electronic Information and Electrical Engineering, Changsha University, Changsha, 410022, China
weishengfei@nudt.edu.cn, suyuanliu@nudt.edu.cn, wang_jun@nudt.edu.cn, liangke200694@126.com,
miaomiaolinudt@gmail.com, l.luo@nudt.edu.cn

Abstract

Fair clustering is crucial for mitigating bias in unsupervised learning, yet existing algorithms often suffer from quadratic or super-quadratic computational complexity, rendering them impractical for large-scale datasets. To bridge this gap, we introduce the Anchor-based Fair Clustering Framework (AFCF), a novel, general, and plug-and-play framework that empowers arbitrary fair clustering algorithms with linear-time scalability. Our approach first selects a small but representative set of anchors using a novel fair sampling strategy. Then, any off-the-shelf fair clustering algorithm can be applied to this small anchor set. The core of our framework lies in a novel anchor graph construction module, where we formulate an optimization problem to propagate labels while preserving fairness. This is achieved through a carefully designed group-label joint constraint, which we prove theoretically ensures that the fairness of the final clustering on the entire dataset matches that of the anchor clustering. We solve this optimization efficiently using an ADMM-based algorithm. Extensive experiments on multiple large-scale benchmarks demonstrate that AFCF drastically accelerates state-of-the-art methods, which reduces computational time by orders of magnitude while maintaining strong clustering performance and fairness guarantees.

Code — https://github.com/smcsurvey/AFCF
Extended version — https://arxiv.org/abs/2511.09889

Introduction

Machine learning has been widely applied in key domains such as finance, education, and healthcare. It is concerning that models(Zhou et al. 2025; Wang et al. 2024a; Yang et al. 2025b) may exhibit discrimination against groups characterized by sensitive attributes such as race and gender, due to biases inherent in the data(Chouldechova and Roth 2020; Buolamwini and Gebru 2018). Recently, such issues have sparked extensive interest within the community regarding algorithmic fairness in supervised learning(Zafar et al. 2017; Donini et al. 2018; Huang et al. 2022) Moreover, the incorporation of algorithmic fairness constraints has begun to be explored in unsupervised learning(Chierichetti et al. 2017; Kleindessner et al. 2019; Li et al. 2024). Specifically, Chierichetti et al.(Chierichetti et al. 2017) pioneered the concept of fair clustering, advocating for approximate proportional representation of samples from each sensitive group within every cluster.

Recent research has focused on develoxping methods to ensure group fairness in clustering models. Several works have specifically addressed fairness in prototype-based clustering, e.g., (Chierichetti et al. 2017; Carreira-Perpinán and Wang 2013; Backurs et al. 2019; Bera et al. 2019). Very recently, Kleindessner et al.(Kleindessner et al. 2019) formalized this notion of group fairness within the spectral clustering framework(Shi and Malik 2000; Von Luxburg 2007). However, constrained by computational complexity and other factors, many existing approaches lack scalability. The pioneering fair clustering approach introduced by Chierichetti et al.(Chierichetti et al. 2017) exhibits super-quadratic runtime complexity, primarily due to its initial fairlet decomposition phase, resulting in significant scalability limitations(Chhabra, Masalkovaitė, and Mohapatra 2021). Similarly, Kleindessner et al.(Kleindessner et al. 2019) incorporated linear fairness constraints on the assignment matrix within a spectral relaxation framework. However, their method necessitates storing the full affinity matrix and computing its eigenvalue decomposition. This incurs cubic complexity for a direct implementation, and while optimized approximations exist, they typically achieve only super-quadratic complexity(Tian et al. 2014), thereby imposing substantial scalability constraints.

Despite progress in enhancing the scalability of fair clustering, significant limitations persist. For instance, Backurs et al.(Backurs et al. 2019) introduced a tree-based metric approach to construct fairlets in near-linear time. However, this method is restricted to settings with only two protected groups. Although Wang et al.(Wang et al. 2023) accelerated computations for fair spectral clustering (SC), the persistent reliance on the computationally expensive eigendecomposition of the fairness-constrained graph Laplacian remains a fundamental bottleneck, limiting its practical adoption. Concurrently, Chhabra et al.(Chhabra et al. 2022) proposed a more general framework for fair clustering using antidote data. However, the authors explicitly note its primary limitation: prohibitively expensive computational costs when ap-

*These authors contributed equally.
†Corresponding authors

plied to high-dimensional or large-scale datasets.

To address these issues, inspired by prototype learning(Qin, Feng, and Zhang 2025), we propose a general anchor-based fair clustering framework (Figure. 1) consists of four modules: Fair Anchor Generation Module, Anchor Fair Clustering Module, Fair Anchor Graph Construction Module, Label Propagation Module. This method introduces anchors into fair clustering to reduce the clustering scope from n to m, where n and m denote the dataset size and number of anchors respectively ($m \ll n$). Specifically, we first select demographically balanced anchors via quota constraints. We then apply any fair clustering algorithm to these anchors to generate cluster labels (supporting plug-and-play algorithm substitution). Next, we solve the fairness-constrained anchor graph matrix using ADMM optimization. Finally, we propagate anchor labels to the entire dataset through label propagation.

This framework confines high-complexity fair clustering operations to the anchor subset, while enabling global propagation of fairness constraints via the anchor graph, allowing computationally intensive fair clustering algorithms to scale efficiently to large-scale datasets.The key contributions of our work can be summarized as follows:

- To address the pervasive quadratic or super-quadratic complexity limitation in existing fair clustering methods, we propose a novel general anchor-based fair clustering framework. This achieves linear-time scalability for arbitrary fair clustering algorithms by reducing the problem to optimizing over m anchors ($m \ll n$), maintaining clustering quality and fairness while lowering complexity from $\mathcal{O}(n^k)$ ($k \geq 2$) to $\mathcal{O}(n)$, providing a universal solution for large-scale fair clustering.
- We introduce a protected group-label co-constraint mechanism during anchor graph construction, which establishes theoretical guarantees for fairness equivalence transfer between anchor-based clustering and global clustering. This fairness-preserving module ensures that non-anchor assignments simultaneously satisfy cluster cohesion and demographic parity through ADMM optimization.
- Extensive experiments with multiple fair clustering algorithms across diverse datasets demonstrate the framework's efficacy. It achieves significant speedup over existing methods while preserving clustering performance and group balance, particularly under increasing data scales.

Related Work

Fair Clustering

In recent years, fair clustering has received growing attention in the machine learning community. Since conventional clustering methods may produce biased outcomes due to the influence of sensitive attributes, significant research efforts have been dedicated to developing fairness-constrained clustering methodologies(Ahmadian et al. 2019; Chhabra, Masalkovaitė, and Mohapatra 2021; Shaham et al. 2025). Chierichetti et al.(Chierichetti et al. 2017) pioneered the

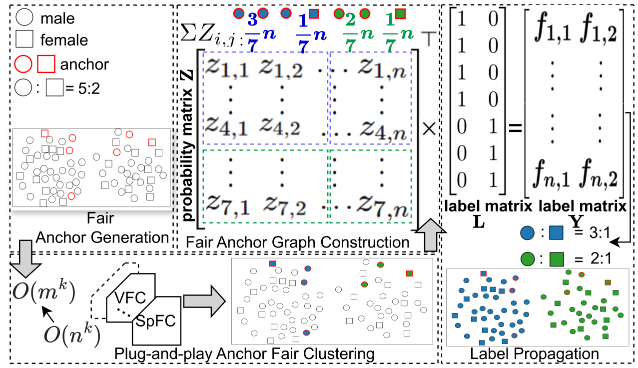

Figure 1: Conceptual framework of the proposed fair anchor-based clustering. Anchors are selected proportionally to cluster demographics, with the left yielding 3 circle and 1 square anchors and the right yielding 2 circle and 1 square anchors. This enables efficient fair clustering on anchor points only, where m $\ll$ n. Group-label joint constraints in the anchor graph $\mathbf{Z}$ maintain demographic proportions with sums of 3/7 for blue circles, 1/7 for blue squares, 2/7 for orange circles, and 1/7 for orange squares. These fairness properties propagate to final clusters through $\mathbf{Y} = \mathbf{Z}^\top \mathbf{L}$, preserving the original demographic ratios. The approach integrates proportional representation, plug-and-play algorithmic flexibility, and constrained graph optimization for fairness preservation.

concept of fair clustering for binary protected groups, while Bera et al.(Bera et al. 2019) extended this framework to multidimensional protected groups, formalizing fairness evaluation through cluster balance metrics.Fair clustering methods quantify fairness through balance metrics defined on protected groups. Specifically, consider a dataset $\mathbf{X} \in \mathbb{R}^{d \times n}$ with n samples and d features, partitioned into c disjoint clusters $\mathcal{C} = \{C_1, C_2, \ldots, C_c\}$. For t mutually exclusive protected groups $\mathcal{G} = \{G_1, G_2, \ldots, G_t\}$, define the global proportion $\rho_r = |G_r|/n$ and intra-cluster proportion $\rho_r^{(l)} = |C_l \cap G_r|/|C_l|$. The cluster fairness metric is formulated as:

$$\text{balance}(\mathcal{C}) = \min_{l \in [c]} \min_{\substack{r, r' \in [t] \\ r \neq r'}} \frac{\rho_r^{(l)}}{\rho_{r'}^{(l)}} \tag{1}$$

where $[c] = \{1, 2, \ldots, c\}$, $[t] = \{1, 2, \ldots, t\}$,with higher values indicating fairer cluster distributions.

Building upon this concept, researchers have proposed diverse approaches to fair clustering from various perspectives. Numerous methods incorporate fairness constraints into the optimization objective. For instance, SpFC(Kleindessner et al. 2019) embeds fairness as linear constraints within spectral clustering. However, due to the necessity of storing the $n \times n$ affinity matrix and computing its eigendecomposition—which remains super-quadratic even for fast implementations—the approach(Tian et al. 2014) suffers from scalability limitations. While Wang et al.(Wang et al. 2023) accelerated computations for fair spectral clustering, their method still relies on the computationally expensive eigendecomposition of the fairness-constrained graph Laplacian. FairSC(Tonin et al. 2025) reformulates the fair spectral clustering problem within a difference-of-convex (DC) programming framework, yet maintains quadratic complexity $\mathcal{O}(n^2)$. VFC(Ziko et al. 2021) incorporates fairness penalties into the clustering objective through a variational framework, achieving scalability on large-scale datasets via decoupled updates of assignment variables.Some approaches employ pre-processing or post-processing techniques. For instance, Chierichetti et al.(Chierichetti et al. 2017) transform raw data into fairlet representations satisfying fairness constraints, enabling fair clustering through classical algorithms. However, this method exhibits quadratic time complexity with respect to the number of data points.Backurs et al.(Backurs et al. 2019) accelerated fairlet decomposition to near-linear time complexity, though their approach remains restricted to binary protected groups. Concurrently,Chhabra et al.(Chhabra, Singla, and Mohapatra 2022) proposed a more general antidote data fair clustering approach, but the authors explicitly note its computational inefficiency on large-scale datasets as a primary limitation.CFC(Chhabra et al. 2022) achieves fairness through data pre-processing or re-sampling techniques. In contrast, AFC(Bera et al. 2019) transforms clustering outputs into fair solutions via linear programming formulations. Additionally, deep fair clustering methods have emerged, exemplified by DFC(Li, Zhao, and Liu 2020) which enforces fairness through adversarial training.

Anchor-Based Clustering

To address the computational challenges of traditional clustering methods(Wang et al. 2024b; Yang et al. 2025a) on large-scale datasets, recent works have introduced anchor-based techniques (Qiang et al. 2021; Nie et al. 2023; Liu et al. 2024) to accelerate and optimize the clustering process.The core principle of anchor-based clustering is to generate a set of representative anchors from the data pool, where the number of anchors m is typically orders of magnitude smaller than the total number of samples(Zhang, Zhang, and Sun 2023; Liu et al. 2022). A similarity graph is then constructed to measure affinities between samples and anchors. Subsequently, clustering procedures are performed on the representation matrix. For example, Wang et al.(Wang et al. 2021) proposed a self-supervised clustering model to learn the similarity matrix; Liu et al.(Liu et al. 2024) developed an iteratively optimized anchor strategy for multi-view clustering tasks; Nie et al.(Nie et al. 2023) established a unified framework based on anchors that accelerates clustering by updating compact label matrices. These anchor-based methods achieve speedup without sacrificing performance when appropriate anchors are selected.

Method

In this section, we present a general anchor-based fair clustering framework for large-scale fair clustering. We first introduce the four modules of this framework in sequence: the Fair Anchor Generation Module, the Anchor Fair Clustering Module, the Fair Anchor Graph Construction Module, and the Label Propagation Module. Finally, we conduct a complexity analysis of the proposed algorithm.

Overview

Existing fair clustering methods face difficulties in scaling to large-scale data due to super-quadratic complexity. To address this, we propose a general Anchor-based Fair Clustering Framework (AFCF) that enables linear-time scalability for arbitrary fair clustering algorithms. First, we employ a novel fair sampling strategy to select a small yet representative anchor set. Then, any off-the-shelf fair clustering algorithm is applied to this anchor subset. Next, through a carefully designed group-label co-constraint, we formulate an optimization problem that ensures fairness during anchor graph construction, with theoretical guarantees that this constraint preserves identical fairness properties between anchor clustering and final full-dataset clustering. Finally, we solve this optimization via ADMM and propagate anchor labels to the entire dataset to obtain the final clustering.

Fair Anchor Generation

Two prevalent anchor selection strategies are random sampling and k-means (Jain 2010). However, the stochastic nature of the former yields unstable results with no performance guarantees, while the latter's sensitivity to initial centroids causes algorithmic instability, often requiring multiple runs to mitigate randomness effects. Crucially, neither approach considers fair representation across protected demographic groups.

Li et al.(Li et al. 2020) proposed a simple yet effective anchor selection strategy called Directly Alternate Sampling (DAS), which selects anchors that comprehensively cover the entire data point cloud. Building on this coverage principle while ensuring fair representation across protected demographic groups, we propose the Fair Directly Alternate Sampling (FDAS) method. Algorithm 1 summarizes the FDAS procedure. FDAS introduces a dual fairness mechanism on top of DAS: first, proportional quota allocation (Step 2) ensures anchor counts per protected group match global demographic proportions by computing base quotas, and distributes the residual anchors Δ to the groups with the smallest current counts to avoid certain groups being completely ignored; second, within-group selection employs nonlinear score decay (Step 15) to promote spatial diversity by suppressing scores near selected samples. This design preserves DAS's spatial coverage capability while significantly improving protected group representation.

Anchor Fair Clustering Module

In the Anchor Fair Clustering module, we apply an arbitrary fair clustering operator to the anchors obtained from FDAS:

$$\mathbf{l} = \mathcal{F}(\mathcal{A}, \mathbf{g}_{\mathcal{A}}, k) \tag{2}$$

where $\mathcal{A}$ denotes the $m \times d$ anchor matrix, $\mathbf{g}_{\mathcal{A}}$ represents the sensitive attribute vector for anchors, k specifies the number of clusters, $\mathcal{F}$ corresponds to the fair clustering operator interface, and $\mathbf{l}$ is the $m \times 1$ cluster assignment vector. Given

Algorithm 1: FDAS

Require: Data matrix $\mathbf{X} \in \mathbb{R}^{d \times n}$, Number of anchors m, Group proportion vector $\mathbf{t} \in \mathbb{R}^{t}$, Group labels group $\in \mathbb{Z}^{n}$
Ensure: Anchor indices $\mathcal{A}$

```
Normalize data: X ← X − min(X)
Compute group quotas: counts ← ⌊m · t⌋
Δ ← m − Σ_{g=0}^{t−1} counts[g]
for k = 1 to Δ do
    g* ← arg min_g counts[g]
    counts[g*] ← counts[g*] + 1
end for
for g = 0 to t − 1 do
    G ← {i | group[i] = g}
    s ← [ Σ_{p=1}^{d} X_{p,i} ]_{i∈G}
    s ← s / max(s)
    for j = 1 to counts[g] do
        k ← arg max s
        A ← A ∪ {G[k]}
        s ← s ⊙ (1 − s) / max(s) {⊙ denotes element-wise multiplication}
    end for
end for
return A
```

$m \ll n$, this module enables seamless substitution with various fair clustering algorithms $\mathcal{F} \in \{\text{SPFC}, \text{VFC}, \cdots\}$, effectively circumventing scalability bottlenecks in large-scale clustering scenarios.

Fair Anchor Graph Construction Module

Problem Formulation To ensure lossless transfer of fairness from anchor clustering to the final partitioning, we propose a provable fairness maintenance mechanism (Proposition 1). Unlike traditional anchor graph methods that solely optimize reconstruction accuracy, our approach introduces intersectional group-label constraints that preserve demographic parity across sensitive attributes and cluster assignments.

Formally, given data matrix $\mathbf{X} \in \mathbb{R}^{d \times n}$ and anchor features $\mathbf{H} \in \mathbb{R}^{d \times m}$, we construct the anchor graph $\mathbf{Z} \in \mathbb{R}^{m \times n}$ by solving:

$$\min_{\mathbf{Z}} \|\mathbf{X} - \mathbf{H}\mathbf{Z}\|_F^2 + \alpha \|\mathbf{Z}\|_F^2 \tag{3}$$

$$s.t. \sum_{j \in \mathcal{C}_l} \sum_{i \in G_r} \mathbf{Z}_{j,i} = t_{l,r}, \quad \forall l \in [c], r \in [t],$$
$$\mathbf{Z}_{:,i} \in \Delta^m, \quad \forall i$$

where $\Delta^m := \{z \in \mathbb{R}^m \mid z_j \geq 0, \sum_j z_j = 1\}$ denotes the standard simplex, $\mathcal{G}_{l,r}$ denotes anchor subgroups defined by the conjunction of sensitive attribute r and cluster label l, and $t_{l,r} = |\mathcal{G}_{l,r}| \cdot n/m$. This constraint enforces proportional representation consistency, as formally stated in Proposition 1, whose proof is provided in the appendix.

Proposition 1. *Under the fairness constraints in (3), the balance metric balance($\mathcal{C}$) of the final clustering equals that of the anchor clustering balance($\mathcal{C}_a$):*

$$\textit{balance}(\mathcal{C}) = \textit{balance}(\mathcal{C}_a) \tag{4}$$

ADMM Optimization To efficiently solve the constrained optimization problem in Eq. (3), we introduce an auxiliary variable $\mathbf{E}$ with equality constraint $\mathbf{Z} = \mathbf{E}$. This strategic decoupling separates the simplex constraints from the fairness conditions, isolates the fairness constraints to the $\mathbf{E}$-subproblem while maintaining the probability simplex constraints on $\mathbf{Z}$, connected through the consensus constraint $\mathbf{Z} = \mathbf{E}$ (Boyd et al. 2011).

We adopt the Alternating Direction Method of Multipliers (ADMM) to solve. Specifically, the augmented Lagrangian is:

$$\begin{aligned}\mathcal{L}_\rho(\mathbf{Z}, \mathbf{E}, \mathbf{\Lambda}) =& \|\mathbf{X} - \mathbf{H}\mathbf{Z}\|_F^2 + \alpha \|\mathbf{Z}\|_F^2 + \langle \mathbf{\Lambda}, \mathbf{Z} - \mathbf{E} \rangle \\ &+ \frac{\rho}{2} \|\mathbf{Z} - \mathbf{E}\|_F^2\end{aligned} \tag{5}$$

where $\mathbf{\Lambda} \in \mathbb{R}^{m \times n}$ is the dual variable and $\rho > 0$ is the penalty parameter. The ADMM iterations proceed as:

$$\mathbf{Z}^{k+1} = \underset{\mathbf{Z} \in \Delta^m}{argmin}\, \mathcal{L}_\rho(\mathbf{Z}, \mathbf{E}^k, \mathbf{\Lambda}^k) \tag{6}$$

$$\mathbf{E}^{k+1} = \underset{E}{argmin}\, \mathcal{L}_\rho(\mathbf{Z}^{k+1}, \mathbf{E}, \mathbf{\Lambda}^k) \tag{7}$$

$$\mathbf{\Lambda}^{k+1} = \mathbf{\Lambda}^k + \rho(\mathbf{Z}^{k+1} - \mathbf{E}^{k+1}) \tag{8}$$

Our algorithm is summarized in Algorithm 2. The penalty $\rho^{(i+1)}$ is updated according to the standard rule suggested in (Boyd et al. 2011) and detailed in Appendix. We now analyze the two subproblems separately.

Solution to Z-subproblem The $\mathbf{Z}$-update in Eq. (6) decomposes into n independent subproblems along data points:

$$\min_{z_i \in \Delta^m} \frac{1}{2} z_i^\top \mathbf{Q} z_i + c_i^\top z_i \tag{9}$$

with $\mathbf{Q} = 2(\mathbf{H}^\top \mathbf{H} + (\alpha + \rho/2)\mathbf{I})$ and $c_i = -2\mathbf{H}^\top x_i - \rho e_i^k + \lambda_i^k$. We solve this quadratic program over the simplex using the Frank-Wolfe algorithm (Ñanculef et al. 2014), selected for its projection-free property and linear convergence rate in convex problems over polytopes. The complete procedure is detailed in Appendix.

Solution to E-subproblem The $\mathbf{E}$-update in Eq. (7) requires projection onto the fairness constraints. Defining $\mathbf{R} = \mathbf{Z}^{k+1} + \rho^{-1}\mathbf{\Lambda}^k$, we obtain the closed-form solution:

$$\mathbf{E}_{j,i} = \mathbf{R}_{j,i} + \frac{t_{l,r} - \sum_{a \in \mathcal{C}_l} \sum_{b \in G_r} \mathbf{R}_{a,b}}{|\mathcal{C}_l| \cdot |G_r|}, \quad \forall j \in \mathcal{C}_l, i \in G_r \tag{10}$$

where $\mathcal{C}_l$ denotes the set of anchors assigned to cluster l, and G_r denotes the set of samples belonging to sensitive group r. This operation redistributes mass uniformly within each group $\mathcal{C}_l \times G_r$ to satisfy the fairness constraint $\sum_{j \in \mathcal{C}_l} \sum_{i \in G_r} \mathbf{E}_{j,i} = t_{l,r}$.

Algorithm 2: Fair Anchor Graph Construction

Require: Data matrix $\mathbf{X} \in \mathbb{R}^{d\times n}$, Anchor features $\mathbf{H} \in \mathbb{R}^{d\times m}$, Group definitions $\{\mathcal{G}_{l,r}\}$, Regularization parameter $\alpha > 0$, Initial penalty $\rho_0 > 0$, Tolerance $\epsilon > 0$, Maximum iterations K
Ensure: Fair anchor graph $\mathbf{Z} \in \mathbb{R}^{m\times n}$

```
Initialize Z^0 ← 1_m 1_n^⊤ / m
Initialize E^0 ← Z^0, Λ^0 ← 0_{m×n}, k ← 0
repeat
    Update Z^{k+1} according to (6)
    Update E^{k+1} according to (7)
    Update Λ^{k+1} according to (8)
    if k mod 10 = 0 then
        Update ρ {details in Appendix}
    end if
    Compute r_k ← ‖Z^{k+1} − E^{k+1}‖_F
    Compute s_k ← ρ‖E^{k+1} − E^k‖_F
    k ← k + 1
until k > K or max(r_k, s_k) < ε
return Z^k
```

Label Propagation Module

Since the anchor graph $\mathbf{Z}$ satisfies group ratio consistency through the group-label joint constraint in Eq. (3), we propagate anchor labels to the entire dataset via single-step matrix multiplication.This design inherently preserves the fairness property, which grounded in the principle that spatially proximate samples in the feature space exhibit higher likelihood of sharing identical class labels (Xu, Liu, and Geng 2019). Formally, based on the anchor fair clustering results $\mathbf{l}$, we construct the anchor label matrix $\mathbf{L} \in \{0,1\}^{m\times k}$. The $\mathbf{L}$ provides one-hot encoded cluster assignments, where $\mathbf{L}_{i,j} = 1$ iff anchor j belongs to cluster i. Subsequently, label diffusion is achieved through matrix multiplication:

$$\mathbf{Y} = \mathbf{Z}^{\top}\mathbf{L} \tag{11}$$

where $\mathbf{Y} \in \mathbb{R}^{n\times k}$ denotes the probabilistic label assignment matrix, with element $\mathbf{Y}_{ij}$ quantifying the membership likelihood of sample i belonging to class j (Xu, Liu, and Geng 2019). This operation embodies the neighborhood consensus principle in fair clustering: each sample's label distribution emerges as a convex combination of its associated anchors' labels, weighted by their affinity measures $\mathbf{Z}$ (Raghavan, Albert, and Kumara 2007). The discrete cluster assignments $\hat{y}_i$ are then determined via maximum likelihood decision:

$$\hat{y}_i = argmax_{1\leq j\leq k}\mathbf{Y}_{ij}, \quad \forall i \in \{1,\ldots,n\} \tag{12}$$

This module achieves end-to-end label propagation with linear time complexity, eliminating requirements for iterative optimization or post-processing while preserving *fairness-awareness* through geometrically consistent label diffusion.

Complexity Analysis

The AFCF framework achieves linear scalability in sample size n with an overall complexity of $O(nmd+nm^2+nmk)$. This derives from four main components: fair anchor generation requires $O(nd + mn)$ operations for data processing and selection; anchor clustering incurs $O(f(m))$ complexity where f is the embedded algorithm's complexity; fair anchor graph construction costs $O(nm^2 + nmd)$; and label propagation takes $O(nmk)$.

Dataset	Samples	Clusters	Protected Groups
Law School	18,692	2	Gender (2)
Credit	29,537	5	Gender (2)
Bank	41,108	2	Marital status (3)
Zafar	100,000	2	Binary (2)
Census II	2,458,285	5	Gender (2)

Table 1: Datasets used in our experiments.

Experiments

Experiment Setup

Datasets We conduct experiments on five real-world and synthetic fair datasets, including Bank(Moro, Rita, and Cortez 2012), Credit Card(Yeh and Lien 2009), Zafar(Zafar et al. 2017), Law School(Le Quy et al. 2022), and Census II(Cohen-Addad et al. 2025).These datasets are widely used in the field of fair clustering(Tonin et al. 2025; Ziko et al. 2021). Dataset specifications are detailed in Table 1.

Implementation details We incorporate five state-of-the-art fair clustering algorithms into our framework and conduct comparative analysis against their original implementations: spFC (Kleindessner et al. 2019), VFC (Ziko et al. 2021), FFC (Pan and Zhong 2023), FMSC (Li et al. 2024) (adapted from multi-view to single-view operation), and fairletFC (Chierichetti et al. 2017). Hyperparameters for both standalone executions and second-module executions within our framework were configured following recommendations in respective source publications. For the proposed AFCF framework, we perform grid search over two key hyperparameters: the anchor size m is selected from $\{2r, 4r, \ldots, 20r\}$ where r denotes sensitive attribute cardinality, and the balance coefficient α is chosen from $\{0.0001, 0.01, 1, 100\}$. The number of clusters is fixed to the ground-truth class count across all datasets. More details can be found in the appendix. All experiments are conducted on a machine with an AMD Ryzen 7 4800H CPU (8 cores, 2.90 GHz), 32 GB of RAM, and integrated graphics.

Metrics We evaluate clustering performance and fairness using four widely-adopted metrics, all of which follow the "higher-is-better" principle.Clustering performance is evaluated using Accuracy (ACC) and Normalized Mutual Information (NMI), while fairness is quantified through Balance and Minimal Normalized Conditional Entropy (MNCE, quantifying distributional consistency between clusters and global data) (Li et al. 2024), with formal definitions provided below:

$$\text{MNCE} = \frac{\min_{k\in[c]}\left(-\sum_{i=1}^{t}\rho_i^{(k)}\log\rho_i^{(k)}\right)}{-\sum_{i=1}^{t}\rho_i\log\rho_i} \in [0,1] \tag{13}$$

Method	Law School				Credit				Bank				Zafar				Census II			
	ACC	NMI	Balance	MNCE	ACC	NMI	Balance	MNCE	ACC	NMI	Balance	MNCE	ACC	NMI	Balance	MNCE	ACC	NMI	Balance	MNCE
VFC	0.588	0.063	0.762	0.999	0.340	0.150	0.619	0.992	0.642	0.044	0.179	0.972	0.961	0.764	0.652	0.992	0.407	0.243	0.678	0.974
VFC-AF	**0.693**	**0.065**	0.718	0.992	**0.473**	0.124	0.562	0.974	**0.719**	**0.062**	**0.184**	**0.976**	**0.980**	**0.865**	**0.655**	**0.993**	0.329	0.093	0.580	0.949
SpFC	0.852	0.005	0.765	0.999	0.419	0.155	0.570	0.977	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN
SpFC-AF	0.703	**0.074**	**0.765**	**0.999**	**0.419**	0.112	**0.573**	**0.978**	**0.876**	**0.143**	**0.185**	**0.993**	**0.993**	**0.943**	**0.653**	**0.993**	**0.456**	**0.131**	**0.906**	**0.999**
FMSC	0.558	0.060	0.772	1.000	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN
FMSC-AF	**0.651**	**0.067**	0.738	0.996	**0.466**	**0.016**	**0.640**	**0.997**	**0.852**	**0.148**	**0.185**	**0.996**	**0.753**	**0.193**	**0.666**	**0.996**	**0.436**	**0.037**	**0.835**	**0.995**
TFC	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	-	-	-	-	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN
TFC-AF	**0.786**	**0.083**	**0.726**	**0.994**	**0.350**	**0.111**	**0.542**	**0.967**	-	-	-	-	**0.533**	**0.030**	**0.687**	**1.000**	**0.380**	**0.139**	**0.509**	**0.923**
FFC	0.581	0.060	0.758	0.998	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN	NAN
FFC-AF	**0.701**	**0.070**	0.753	**0.998**	**0.359**	**0.062**	**0.584**	**0.982**	**0.719**	**0.062**	**0.184**	**0.976**	**0.617**	**0.044**	**0.685**	**1.000**	**0.385**	**0.104**	**0.393**	**0.859**

Table 2: Comprehensive Method Comparison with AFCF Enhancement. NAN indicates failure to complete within 30 minutes; "-" indicates the method does not support scenarios where the number of sensitive attribute groups exceeds 2.

Method	Law School				Credit				Bank				Zafar				Census II			
	ACC	NMI	Balance	MNCE	ACC	NMI	Balance	MNCE	ACC	NMI	Balance	MNCE	ACC	NMI	Balance	MNCE	ACC	NMI	Balance	MNCE
Random	0.902	0.000	0.772	1.000	0.350	0.076	0.512	0.954	0.716	0.042	0.183	0.977	0.978	0.859	0.653	0.993	0.453	0.184	0.662	0.971
DAS	0.902	0.000	0.772	1.000	0.440	0.105	0.560	0.973	0.721	0.061	0.185	0.977	0.936	0.712	0.650	0.992	0.388	0.045	0.101	0.442
FDAS	0.703	**0.074**	0.765	0.999	0.419	**0.112**	**0.573**	**0.978**	**0.876**	**0.143**	**0.185**	**0.993**	**0.993**	**0.943**	**0.653**	**0.993**	**0.456**	0.131	**0.906**	**0.999**

Table 3: Ablation Study on Anchor Selection Methods. Random: randomly generated anchors; DAS: anchors selected by DAS method; FDAS: anchors selected by our FDAS module.

Method	Datasets				
	Law School	Credit	Bank	Zafar	Census II
VFC	49.316	66.932	50.992	65.135	1494.382
VFC-AF	**21.010**	**23.449**	**47.206**	**46.429**	**918.220**
SpFC	79.367	324.095	NAN	NAN	NAN
SpFC-AF	**8.907**	**14.911**	**34.967**	**66.139**	**731.359**
FMSC	179.795	NAN	NAN	NAN	NAN
FMSC-AF	**13.206**	**9.777**	**12.675**	**40.763**	**713.174**
TFC	NAN	NAN	-	NAN	NAN
TFC-AF	**12.002**	**15.360**	-	**47.512**	**1022.102**
FFC	393.130	NAN	NAN	NAN	NAN
FFC-AF	**11.139**	**17.620**	**25.714**	**55.170**	**1109.128**

Table 4: Execution Time Comparison with AFCF Enhancement (seconds). NAN indicates failure to complete within 30 minutes; "-" indicates the method does not support the scenario.

Experimental Results

Tables 3 and 4 show that integrating the Anchor-based Fair Clustering Framework (AFCF) consistently improves computational efficiency across five algorithms and datasets.

First, the clustering quality metrics (NMI, Balance, ACC, and MNCE) highlight the effectiveness of AFCF in enhancing fairness-aware clustering. For example, when VFC is embedded within the AFCF framework (VFC-AF), the model achieves a slight improvement in NMI across multiple datasets, such as Law School (0.063 to 0.065) and Bank (0.044 to 0.062), while also maintaining or improving accuracy. On the other hand, SpFC-AF shows significant improvements, such as moving NMI from near zero (0.005) to 0.074 on Law School, and most importantly, overcoming computational bottlenecks, completing tasks on datasets (like Bank) where SpFC failed to finish.

In terms of computational efficiency, embedding the algorithms within AFCF significantly reduces execution time. SpFC-AF, which previously failed to run due to time constraints, now completes tasks within time complexity that scales linearly with the data size, demonstrating the efficiency gains of embedding fairness constraints into our framework. FMSC-AF, which initially failed to handle large datasets, now successfully processes them within an acceptable time, with runtimes under 10 seconds for certain datasets like Credit. Notably, the FFC-AF method, which has the longest runtime among the methods tested, still benefits from AFCF's optimization, completing in 1109 seconds on Census II, a marked advantage compared to the 1494 seconds required by VFC alone.

Ablation Studies

Table 3 demonstrates the critical importance of our fairness-aware anchor selection (FDAS) module. The complete clustering failure (NMI=0) observed with both Random and DAS methods on Law School highlights the necessity of fairness constraints in anchor selection for sensitive datasets. FDAS not only resolves this degenerate clustering issue but also achieves significant improvements across multiple

Method	Law School				Credit				Bank				Zafar				Census II			
	ACC	NMI	Balance	MNCE	ACC	NMI	Balance	MNCE	ACC	NMI	Balance	MNCE	ACC	NMI	Balance	MNCE	ACC	NMI	Balance	MNCE
AC	0.665	0.074	0.769	1.000	0.417	0.110	0.572	0.977	0.873	0.136	0.183	0.993	0.992	0.935	0.652	0.992	0.465	0.018	0.461	0.900
FAC	**0.703**	**0.074**	0.765	0.999	**0.419**	**0.112**	**0.573**	**0.978**	**0.876**	**0.143**	**0.185**	**0.993**	**0.993**	**0.943**	**0.653**	**0.993**	0.456	**0.131**	**0.906**	**0.999**

Table 5: Comparison between FAC and AC Methods. FAC: full proposed method; AC: ablated version.

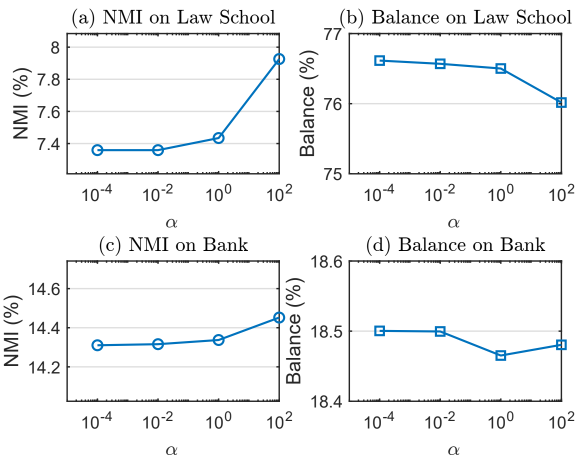

Figure 2: NMI and Balance on Law School and Bank data sets w.r.t. different values of α.

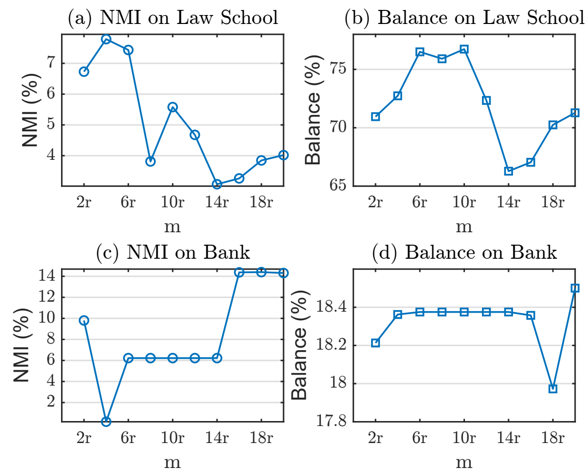

Figure 3: NMI and Balance on Law School and Bank data sets w.r.t. different values of m.

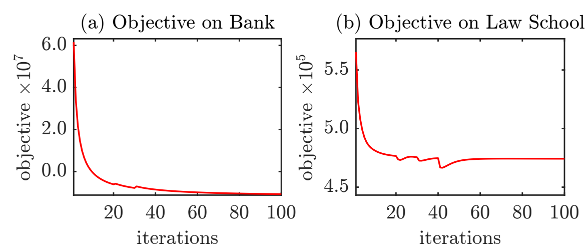

Figure 4: The convergence of the proposed algorithm for minimizing the objective in (5). The plots are based on the Bank and Law School datasets.

benchmarks. On Bank data, FDAS substantially outperforms baselines in both clustering quality and fairness metrics. Most notably, for Census II, FDAS provides dramatic fairness improvements while maintaining competitive accuracy, confirming that FDAS is essential for preventing bias propagation in downstream clustering tasks.

Table 5 validates the necessity of our fairness-aware anchor graph construction (FAC) module. The results demonstrate that explicitly incorporating fairness constraints during graph construction consistently enhances both clustering quality and fairness metrics. Particularly on Census II, FAC significantly improves balance while maintaining competitive accuracy, confirming that fairness constraints prevent biased representation propagation in sensitive datasets. The observed improvements across multiple datasets highlight that fair graph construction is crucial for equitable clustering outcomes.

Sensitivity and Convergence Analysis

The robustness of our method is evaluated through systematic parameter variations. Figure 2 shows NMI mildly increases with α (10^{-4}–10^{2}) while balance remains stable, confirming fairness preservation is α-insensitive.

Figure. 3, Law School peaks at $m = 4r$ (NMI) and $m = 10r$ (balance), whereas Bank requires $m \geq 16r$ for optimal NMI with consistently high balance. Fairness metrics thus exhibit stronger robustness to anchor selection than clustering quality.

Figure 4 shows convergence behavior of Bank and Law School, which illustrates that the objective value of our algorithm consistently decreases with each iteration, which provides clear evidence of the convergence of our proposed algorithm.

Conclusion

This work addresses the critical scalability limitations in fair clustering through the Anchor-based Fair Clustering Framework (AFCF). By introducing a novel fair sampling strategy and fairness-preserving anchor graph construction with group-label joint constraints, AFCF enables linear-time complexity for arbitrary fair clustering algorithms. Our theoretical analysis guarantees global fairness equivalence with anchor-level clustering, while extensive experiments demonstrate orders-of-magnitude acceleration. The framework making large-scale fair clustering practically feasible.

Acknowledgements

This work is supported by the National Natural Science Foundation of China (No.625B2182, 62376039 and 62506371).

References

Ahmadian, S.; Epasto, A.; Kumar, R.; and Mahdian, M. 2019. Clustering without over-representation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 267–275.

Backurs, A.; Indyk, P.; Onak, K.; Schieber, B.; Vakilian, A.; and Wagner, T. 2019. Scalable fair clustering. In *International conference on machine learning*, 405–413. PMLR.

Bera, S.; Chakrabarty, D.; Flores, N.; and Negahbani, M. 2019. Fair algorithms for clustering. *Advances in Neural Information Processing Systems*, 32.

Boyd, S.; Parikh, N.; Chu, E.; Peleato, B.; Eckstein, J.; et al. 2011. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine learning*, 3(1): 1–122.

Buolamwini, J.; and Gebru, T. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, 77–91. PMLR.

Carreira-Perpinán, M. A.; and Wang, W. 2013. The K-modes algorithm for clustering. *arXiv preprint arXiv:1304.6478*.

Chhabra, A.; Li, P.; Mohapatra, P.; and Liu, H. 2022. Robust fair clustering: A novel fairness attack and defense framework. *arXiv preprint arXiv:2210.01953*.

Chhabra, A.; Masalkovaitė, K.; and Mohapatra, P. 2021. An overview of fairness in clustering. *IEEE Access*, 9: 130698–130720.

Chhabra, A.; Singla, A.; and Mohapatra, P. 2022. Fair clustering using antidote data. In *Algorithmic Fairness through the Lens of Causality and Robustness workshop*, 19–39. PMLR.

Chierichetti, F.; Kumar, R.; Lattanzi, S.; and Vassilvitskii, S. 2017. Fair clustering through fairlets. *Advances in neural information processing systems*, 30.

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

Cohen-Addad, V.; Jiang, S. H.-C.; Yang, Q.; Zhang, Y.; and Zhou, S. 2025. Fair Clustering in the Sliding Window Model. *arXiv preprint arXiv:2503.05173*.

Donini, M.; Oneto, L.; Ben-David, S.; Shawe-Taylor, J. S.; and Pontil, M. 2018. Empirical risk minimization under fairness constraints. *Advances in neural information processing systems*, 31.

Huang, X.; Li, Z.; Jin, Y.; and Zhang, W. 2022. Fair-AdaBoost: Extending AdaBoost method to achieve fair classification. *Expert Systems with Applications*, 202: 117240.

Jain, A. K. 2010. Data clustering: 50 years beyond K-means. *Pattern recognition letters*, 31(8): 651–666.

Kleindessner, M.; Samadi, S.; Awasthi, P.; and Morgenstern, J. 2019. Guarantees for spectral clustering with fairness constraints. In *International conference on machine learning*, 3458–3467. PMLR.

Le Quy, T.; Roy, A.; Iosifidis, V.; Zhang, W.; and Ntoutsi, E. 2022. A survey on datasets for fairness-aware machine learning. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 12(3): e1452.

Li, P.; Zhao, H.; and Liu, H. 2020. Deep fair clustering for visual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9070–9079.

Li, R.; Hu, H.; Du, L.; Chen, J.; Jiang, B.; and Zhou, P. 2024. One-stage fair multi-view spectral clustering. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 1407–1416.

Li, X.; Zhang, H.; Wang, R.; and Nie, F. 2020. Multiview clustering: A scalable and parameter-free bipartite graph fusion method. *IEEE transactions on pattern analysis and machine intelligence*, 44(1): 330–344.

Liu, S.; Liang, K.; Dong, Z.; Wang, S.; Yang, X.; Zhou, S.; Zhu, E.; and Liu, X. 2024. Learn from view correlation: An anchor enhancement strategy for multi-view clustering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 26151–26161.

Liu, S.; Liu, X.; Wang, S.; Niu, X.; and Zhu, E. 2022. Fast incomplete multi-view clustering with view-independent anchors. *IEEE Transactions on Neural Networks and Learning Systems*, 35(6): 7740–7751.

Moro, S.; Rita, P.; and Cortez, P. 2012. Bank marketing. *UCI Machine Learning Repository*, 10: C5K306.

Ñanculef, R.; Frandi, E.; Sartori, C.; and Allende, H. 2014. A novel frank–wolfe algorithm. analysis and applications to large-scale svm training. *Information Sciences*, 285: 66–99.

Nie, F.; Xue, J.; Yu, W.; and Li, X. 2023. Fast clustering with anchor guidance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(4): 1898–1912.

Pan, R.; and Zhong, C. 2023. Fairness first clustering: A multi-stage approach for mitigating bias. *Electronics*, 12(13): 2969.

Qiang, Q.; Zhang, B.; Wang, F.; and Nie, F. 2021. Fast multi-view discrete clustering with anchor graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 9360–9367.

Qin, Y.; Feng, G.; and Zhang, X. 2025. Scalable One-Pass Incomplete Multi-View Clustering by Aligning Anchors. In *Advancement of Artificial Intelligence*, 20042–20050.

Raghavan, U. N.; Albert, R.; and Kumara, S. 2007. Near linear time algorithm to detect community structures in large-scale networks. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 76(3): 036106.

Shaham, S.; Hajisafi, A.; Quan, M. K.; Nguyen, D. C.; Krishnamachari, B.; Peris, C.; Ghinita, G.; Shahabi, C.; and Pathirana, P. N. 2025. Privacy and fairness in machine learning: A survey. *IEEE Transactions on Artificial Intelligence*.

Shi, J.; and Malik, J. 2000. Normalized cuts and image segmentation. *IEEE Transactions on pattern analysis and machine intelligence*, 22(8): 888–905.

Tian, F.; Gao, B.; Cui, Q.; Chen, E.; and Liu, T.-Y. 2014. Learning deep representations for graph clustering. In *Proceedings of the AAAI conference on artificial intelligence*, volume 28.

Tonin, F.; Lambert, A.; Suykens, J. A.; and Cevher, V. 2025. Accelerating Spectral Clustering under Fairness Constraints. *arXiv preprint arXiv:2506.08143*.

Von Luxburg, U. 2007. A tutorial on spectral clustering. *Statistics and computing*, 17(4): 395–416.

Wang, F.; Jin, J.; Dong, Z.; Yang, X.; Feng, Y.; Liu, X.; Zhu, X.; Wang, S.; Liu, T.; and Zhu, E. 2024a. View gap matters: Cross-view topology and information decoupling for multi-view clustering. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 8431–8440.

Wang, F.; Jin, J.; Hu, J.; Liu, S.; Yang, X.; Wang, S.; Liu, X.; and Zhu, E. 2024b. Evaluate then cooperate: Shapley-based view cooperation enhancement for multi-view clustering. *Advances in Neural Information Processing Systems*, 37: 135355–135379.

Wang, J.; Lu, D.; Davidson, I.; and Bai, Z. 2023. Scalable spectral clustering with group fairness constraints. In *International conference on artificial intelligence and statistics*, 6613–6629. PMLR.

Wang, J.; Ma, Z.; Nie, F.; and Li, X. 2021. Fast self-supervised clustering with anchor graph. *IEEE Transactions on Neural Networks and Learning Systems*, 33(9): 4199–4212.

Xu, N.; Liu, Y.-P.; and Geng, X. 2019. Label enhancement for label distribution learning. *IEEE Transactions on Knowledge and Data Engineering*, 33(4): 1632–1643.

Yang, X.; Wang, S.; Jin, J.; Wang, F.; Liu, T.; Jin, Y.; Liu, X.; Zhu, E.; and He, K. 2025a. Generalized Deep Multi-view Clustering via Causal Learning with Partially Aligned Cross-view Correspondence. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.

Yang, X.; Wang, S.; Wang, F.; Jin, J.; Liu, S.; Liu, Y.; Zhu, E.; Liu, X.; and Jin, Y. 2025b. Automatically Identify and Rectify: Robust Deep Contrastive Multi-view Clustering in Noisy Scenarios. In *International Conference on Machine Learning*. PMLR.

Yeh, I.-C.; and Lien, C.-h. 2009. The comparisons of data mining techniques for the predictive accuracy of probability of default of credit card clients. *Expert systems with applications*, 36(2): 2473–2480.

Zafar, M. B.; Valera, I.; Rogriguez, M. G.; and Gummadi, K. P. 2017. Fairness constraints: Mechanisms for fair classification. In *Artificial intelligence and statistics*, 962–970. PMLR.

Zhang, N.; Zhang, X.; and Sun, S. 2023. Efficient multi-view representation learning with correntropy and anchor graph. *IEEE Transactions on Knowledge and Data Engineering*, 36(9): 4632–4645.

Zhou, T.; Dong, Z.; Wang, S.; Liang, K.; Li, M.; Liu, X.; Zhu, E.; and Dong, X. 2025. DPFMVC: Dynamic Progressive Fusion for Multi-view Clustering. In *Proceedings of the 33rd ACM International Conference on Multimedia*, MM '25, 1102–1111. New York, NY, USA: Association for Computing Machinery. ISBN 9798400720352.

Ziko, I. M.; Yuan, J.; Granger, E.; and Ayed, I. B. 2021. Variational fair clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 11202–11209.