

# Highly Efficient Rotation-Invariant Spectral Embedding for Scalable Incomplete Multi-View Clustering

Xinxin Wang<sup>1</sup>, Yongshan Zhang<sup>2</sup>, Yicong Zhou<sup>1\*</sup>

<sup>1</sup>Department of Computer and Information Science, University of Macau

<sup>2</sup>School of Computer Science, China University of Geosciences

xinxinwang1024@gmail.com, yszhang.cug@gmail.com, yicongzhou@um.edu.mo

## Abstract

Incomplete multi-view clustering presents significant challenges due to missing views. Although many existing graph-based methods aim to recover missing instances or complete similarity matrices with promising results, they still face several limitations: (1) Recovered data may be unsuitable for spectral clustering, as these methods often ignore guidance from spectral analysis; (2) Complex optimization processes require high computational burden, hindering scalability to large-scale problems; (3) Most methods do not address the rotational mismatch problem in spectral embeddings. To address these issues, we propose a highly efficient rotation-invariant spectral embedding (RISE) method for scalable incomplete multi-view clustering. RISE learns view-specific embeddings from incomplete bipartite graphs to capture the complementary information. Meanwhile, a complete consensus representation with second-order rotation-invariant property is recovered from these incomplete embeddings in a unified model. Moreover, we design a fast alternating optimization algorithm with linear complexity and promising convergence to solve the proposed formulation. Extensive experiments on multiple datasets demonstrate the effectiveness, scalability, and efficiency of RISE compared to the state-of-the-art methods.

## Introduction

Multi-view clustering (MVC) explores the affinity relationships of samples among multiple views and groups similar samples into the same cluster (Zhang et al. 2018). Over the past decade, a variety of MVC methods have been proposed for various real-world applications (Chen et al. 2020; Li et al. 2024a). Although these methods achieve superior clustering performance, most of them strictly rely on the assumption that all instances in each view are available and complete. However, in real-world scenarios, some instances may be inaccessible from certain views due to transmission loss or sensor malfunctions (Wen et al. 2020). For instance, in medical diagnosis, patients may undergo various pathology tests, such as MRI, CT, and genetic testing. Only a small percentage of patients require all these tests, while others may need only a few targeted assessments. The resulting multi-view data, which contains missing instances,

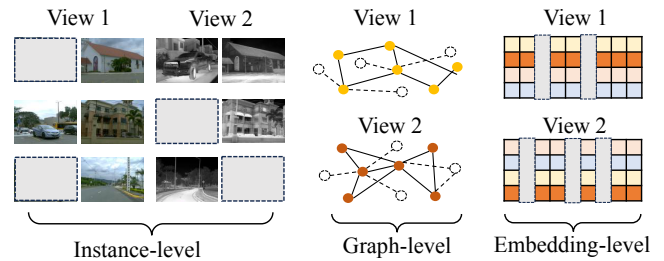


Figure 1: Different completion stages for IMVC. Filling missing data at the feature embedding level is computationally efficient and provides robust performance.

is referred to as incomplete multi-view data. When existing MVC methods are applied to such incomplete datasets, they may experience significant performance degradation or even complete failure (Xia et al. 2022; Shen et al. 2023).

To solve this problem, incomplete multi-view clustering (IMVC) has gained increasing attention. Recently, a variety of IMVC methods have been proposed (Wen, Xu, and Liu 2020; Wen et al. 2021a; Li et al. 2022). These methods can be divided into three categories: instance-level, graph-level, and embedding-level methods. As shown in Fig. 1 (a), instance-level methods are natural and intuitive, as they recover missing instances to facilitate subsequent clustering (Hao et al. 2023; Zhao et al. 2023a; Liu et al. 2021). Typical recovery strategies include simple averaging (Shao, He, and Yu 2015; Gao, Peng, and Jian 2016), matrix factorization (Yin and Sun 2023; Wen et al. 2019), and deep learning-based imputation (Lin et al. 2021). As shown in Fig. 1 (b), graph-level methods aim to fill in incomplete similarity matrices rather than directly recovering missing instances (Wen et al. 2023a; Cui et al. 2022; Liu et al. 2023; Zhang et al. 2023a; Lv et al. 2022; Zhang, Huang, and Wang 2024). During the learning process, the recovered similarity matrices can either be used in spectral clustering or employed to directly output clustering results based on the connected components. Unlike instance-level and graph-level completion methods, as shown in Fig. 1(c), embedding-level methods perform completion in the latent feature space (Zhang et al. 2023b). This approach efficiently recovers missing data from low-dimensional embeddings, and the recovered representations naturally align with clustering tasks.

Although there are viable precedents for filling in miss-

\*Corresponding author.

ing data at different completion stages, existing methods still exhibit several limitations. *First*, most instance-level and graph-level methods finish multi-view fusion and completion before spectral analysis, which may render the recovered data unsuitable for spectral clustering (Zhang et al. 2023b). *Second*, most graph-level methods involve complex optimization processes and require a computational burden of at least  $\mathcal{O}(n^2)$  in both computation and storage, making them impractical for large-scale problems (Liu et al. 2022b; Wang et al. 2022a). *Third*, as is shown in Fig. 2, disparities in available instances among views cause clustering structure shifts across similarity graphs and increase the imbalance of eigenvalue distributions among the Laplacian matrices. This may worsen potential feature rotational mismatches between relaxed spectral embedding matrices when performing SVD decomposition. We refer to this issue as the spectral embedding rotational mismatch (SERM) problem. A similar issue in multi-view subspace learning is referred to as the misalignment problem (Wen et al. 2023b; Wang et al. 2022b). To the best of our knowledge, no unified IMVC framework simultaneously addresses these issues.

To address these issues, this paper proposes a highly efficient rotation-invariant spectral embedding (RISE) method for scalable incomplete multi-view clustering. Specifically, we learn view-specific embeddings from incomplete bipartite graphs to capture discrepancies between views. To tackle the SERM problem among view-specific embeddings, we propose a second-order rotation-invariant learning module to recover the complete consensus representation. Both processes are integrated into a unified framework, where the view-specific embedding can better serve for the learning of complete consensus representation, and the latter is able to guide the learning of the former. To solve the proposed formulation, we propose a fast alternating optimization algorithm with promising convergence, reducing the time complexity from  $\mathcal{O}(n^3v)$  to  $\mathcal{O}(nmv)$ . The storage cost of RISE is approximately  $\mathcal{O}(n)$  by using bipartite graphs. The main contributions of this paper are listed below.

- We propose an efficient IMVC method, termed as RISE. It integrates incomplete view-specific embedding learning and complete consensus representation learning in a unified model. By introducing incomplete bipartite graph, RISE shows scalability and high-efficiency.
- To address the SERM problem, we propose a second-order rotation-invariant learning module to recover the complete consensus representation from view-specific embeddings. Guided by this module, the cross-view spectral embeddings can achieve adequate consistency.
- We design a fast two-step optimization algorithm for RISE with theoretical analyses. Our proposed algorithm demonstrates promising convergence and linear time complexity with respect to the number of samples. Comprehensive experiments validate the effectiveness and efficiency of our method.

## Related Work

To tackle the scalability and efficiency issues in IMVC, several bipartite graph-based approaches have achieved signif-

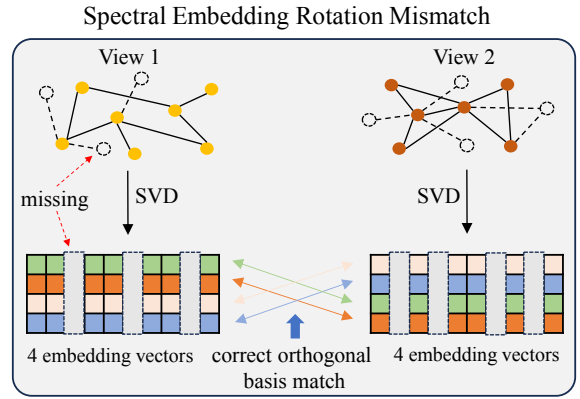


Figure 2: Illustration of the SERM problem. Each color row represents a relaxed orthogonal basis. With structure drift in Laplacian matrices for incomplete views, the resulting orthogonal spectral embedding may be rotational mismatch.

icant attention. Wang *et al.* (Wang et al. 2022a) first introduced the bipartite graph for IMVC. Li *et al.* (Li et al. 2024b) proposed a scalable parameter-free model, while Liu *et al.* (Liu et al. 2022a) developed a fast model with view-independent anchors. Zhao *et al.* (Zhao et al. 2023b) refined the complete bipartite graph learning by using concatenated multiple features. These methods learn a complete unified bipartite graph across views, to which spectral clustering is applied to obtain final clustering results. However, enforcing consistent cross-view representation before spectral analysis can neglect inter-view discrepancies, leading to suboptimal spectral embedding representations. Recent works have integrated tensor learning techniques to explore the inter-view low-rank structure. Chen *et al.* (Chen et al. 2023) combined tensor learning with view-specific bipartite graphs, while Li *et al.* (Li et al. 2023) integrated the consistent distribution between anchors and observed embedded representations to improve anchor imputation. Long *et al.* (Long et al. 2024) applied tensor networks to explore both cross-view and structural correlations. However, computing high-order correlations across views using tensor singular value decomposition (t-SVD) or tensor ring (TR) decomposition can be time-consuming, which limits their applicability in real-world scenarios.

## Proposed Method

### Formulation

Recently, several works have successfully adopted bipartite graphs to enhance the model’s scalability and clustering efficiency in IMVC tasks (Wang et al. 2022a; Li et al. 2024b). Similar to previous works (He et al. 2020), we use bipartite graph to reduce structural redundancy when presenting the correlations among samples. Specifically, given incomplete multi-view data  $\{X^{(i)} \in \mathbb{R}^{d_i \times n_i}\}_{i=1}^v$  with  $d_i$  features and  $n_i$  samples in  $v$  views,  $K$ -means is used to generate  $m$  anchors  $A^{(i)} = [a_1, a_2, \dots, a_m] \in \mathbb{R}^{d_i \times m}$  for each view separately. Then, a bipartite graph  $\overline{B}^{(i)} \in \mathbb{R}^{n_i \times m}$  is constructed between  $m$  anchors and  $n_i$  samples on each view individu-

ally, where each element in  $\overline{B}^{(i)}$  is computed by

$$b_{p,q} = \begin{cases} \frac{d(p,k+1)-d(p,q)}{kd(p,k+1)-\sum_{q=1}^k d(p,q)} & \forall q \in \Phi_p \\ 0 & \forall q \notin \Phi_p \end{cases}, \quad (1)$$

where  $\Phi_p$  is the index set of the  $k$  nearest anchors of sample  $x_p$ , and  $d(p,q) = \|\mathbf{x}_p - \mathbf{a}_q\|^2$  presents the Euclidean distance between  $x_p$  and its  $q$ -th nearest anchor  $a_q$ .

Rather than using full Laplacian matrices, we define a highly efficient spectral embedding learning framework for scalable incomplete multi-view clustering based on bipartite graphs as follows:

$$\max_{F^{(i)}} \sum_{i=1}^v \text{Tr} \left( F^{(i)T} B^{(i)} B^{(i)T} F^{(i)} \right), \text{ s.t. } F^{(i)T} F^{(i)} = I_k, \quad (2)$$

where  $B^{(i)} = \overline{B}^{(i)} \Lambda^{(i)-1/2}$  is the normalized bipartite graph of the  $i$ -th view, and  $F^{(i)} \in \mathbb{R}^{n_i \times k}$  is the corresponding embedding generated from  $B^{(i)}$ . It should be noted that  $\Lambda^{(i)} \in \mathbb{R}^{m \times m}$  is a diagonal matrix with diagonal elements given by  $\Lambda_{q,q}^{(i)} = \sum_{p=1}^{n_i} \overline{B}_{p,q}^{(i)}$ .

By solving Eq. (2), we obtain the embedding  $F^{(i)} \in \mathbb{R}^{n_i \times k}$  with varying numbers of available feature representations. Afterward, a question that naturally arises is: **How to efficiently and effectively refine the cross-view spectral embedding match and recover the consensus complete embedding?** A common solution is to adopt the following formula to learn the complete consensus representation from incomplete view-specific embeddings (Wen et al. 2021a) by

$$\min_Y \sum_{i=1}^v \left\| F^{(i)} - Q^{(i)T} Y \right\|_F^2, \quad (3)$$

where  $Q^{(i)} \in \{0, 1\}^{n \times n_i}$  is an index matrix of the  $i$ -th view that is used to remove the entries corresponding to the unavailable instances from the complete consensus representation  $Y \in \mathbb{R}^{n \times k}$ .  $Q^{(i)}$  is constructed from the index vector  $h^{(i)} \in \mathbb{R}^{n_i}$ , which records the indices of the  $n_i$  instances in  $F^{(i)}$  within the complete dataset. For the  $q$ -th column of  $Q^{(i)}$ , the elements are as follows:

$$Q_{pq}^{(i)} = \begin{cases} 1, & \text{if } p \text{ equals to the } q\text{-th index in } h^{(i)} \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

However, the recovery strategy in Eq. (3) only considers first-order correlations during information fusion and overlooks the potential SERM problem. It may damage the representation capability of the model. Inspired by (Xia, Yang, and Yang 2023), we address this issue by recovering the complete second-order correlations among the embeddings of all samples in a quadratic term form.

The separate processes of spectral embedding learning and complete consensus representation learning may result in suboptimal results. Therefore, we combine them into a

unified formula as follows:

$$\begin{aligned} \min_{Y, F^{(i)}} \sum_{i=1}^v \left\| Y Y^T - Q^{(i)} F^{(i)} F^{(i)T} Q^{(i)T} \right\|_F^2 \\ - \beta \sum_{i=1}^v \text{Tr} \left( F^{(i)T} B^{(i)} B^{(i)T} F^{(i)} \right), \quad (5) \\ \text{s.t. } F^{(i)T} F^{(i)} = I, \quad Y^T Y = I, \end{aligned}$$

where the first term indicates the complete consensus representation learning with second-order rotation-invariant property, and the second term represents view-specific spectral embedding learning. Specifically, we impose an orthogonality constraint on  $Y$  to enhance its discriminative power.  $\beta$  is the penalty parameter, and  $Q^{(i)} F^{(i)}$  denotes the recovery of the original arrangement. It is easy to prove that the formulation in Eq. (5) is rotation-invariant (Von Luxburg 2007). This implies that for a spectral embedding  $F_*^{(i)}$ , given any orthogonal rotation matrix  $R$ ,  $F_*^{(i)} R$  is a feasible spectral embedding set for Eq. (5). Thus, RISE is unaffected by potential rotations and always obtains the optimal solution  $\hat{Y}$ .

Although bipartite graphs can reduce space complexity, simply introducing the index matrix  $Q^{(i)} \in \mathbb{R}^{n \times n_i}$  incurs a space burden of  $\mathcal{O}(n^2)$  and computational overhead of  $\mathcal{O}(n^3)$ , which can become a serious issue for large-scale data. To address this problem, we innovatively define two operations for large-scale IMVC issues:  $Q^{(i)} F^{(i)} = \zeta^{-1}(F^{(i)}, h^{(i)})$ ,  $Q^{(i)T} Y = \zeta(Y, h^{(i)})$ , where  $\zeta(\cdot)$  and  $\zeta^{-1}(\cdot)$  denote the selection of instances from the complete representation matrix according to the index vector and its inverse operation, respectively. With this transformation, the space complexity drops from  $\mathcal{O}(vn^2)$  to  $\mathcal{O}(vn)$ .

## Optimization

To solve the optimization problem in Eq. (5), we devise a fast alternating optimization algorithm to update each variable while keeping the other variables fixed.

**Updating  $F^{(i)}$ :** When  $Y$  is fixed, the optimization problem *w.r.t*  $F^{(i)}$  is

$$\max_{F^{(i)T} F^{(i)} = I} \text{Tr} \left( 2F^{(i)T} \hat{Y} \hat{Y}^T F^{(i)} \right) + \beta \text{Tr} \left( F^{(i)T} B^{(i)} B^{(i)T} F^{(i)} \right), \quad (6)$$

where  $\hat{Y} = Q^{(i)T} Y$ . The solution to Eq. (6) is the  $k$  eigenvectors corresponding to the largest  $k$  eigenvalues of  $S^{(i)} = 2\hat{Y} \hat{Y}^T + \beta B^{(i)} B^{(i)T}$ . However, performing eigen decomposition requires at least  $\mathcal{O}(n_i^2 k)$  time complexity, which makes it impractical for large-scale problems.

In this work, we present an alternative approach to efficiently compute the  $k$  eigenvectors of matrix  $S^{(i)}$  based on proposition 1.

**Proposition 1.** Given a set of similarity matrices  $\{S_t\}_{t=1}^{\pi} \in \mathbb{R}^{n \times n}$ , where each matrix satisfies  $S_t = Z_t Z_t^T$ . Defining the concatenated matrix  $Z_{cat} = [Z_1, \dots, Z_t, \dots, Z_{\pi}] \in \mathbb{R}^{n \times (m_1 + m_2 + \dots + m_{\pi})}$  and supposing the singular value decomposition (SVD) of  $Z_{cat}$

---

**Algorithm 1: RISE**


---

**Input:** Normalized bipartite graph  $\{B^{(i)}\}_{i=1}^v$ , available index  $\{h^{(i)}\}_{i=1}^v$ , number of cluster  $c$ , and embedding dimension  $k$ .

**Initialize:** Initialize  $F^{(i)}$  by performing *SVD* on  $B^{(i)}$ .

- 1: **while** not converge **do**
  - 2:   Update  $Y$  by solving Eq. (8);
  - 3:   Update  $\{F^{(i)}\}_{i=1}^v$  by solving Eq. (6);
  - 4: **end while**
  - 5: **Output:** perform  $k$ -means on  $Y$  to achieve the final clustering indicator.
- 

is  $U\Sigma V^T$  (with  $UU^T = I, V^T V = I$ ), we have that:

$$\max_{F^T F = I} \operatorname{Tr} \left( F^T \left( \sum_{t=1}^{\pi} S_t \right) F \right) \Leftrightarrow \min_{F^T F = I} \|Z_{cat} - F F^T Z_{cat}\|_F^2. \quad (7)$$

In Eq. (7), the optimal solution of  $F$  is equal to  $U$ .

**Proof.** Based on the right-hand side of Eq. (7), we can establish the following equivalences:

$$\begin{aligned} \min_{F^T F = I} \|Z_{cat} - F F^T Z_{cat}\|_F^2 &\Leftrightarrow \max_{F^T F = I} \operatorname{Tr} (F^T Z_{cat} Z_{cat}^T F) \\ &\Leftrightarrow \max_{F^T F = I} \operatorname{Tr} \left( F^T \left( \sum_{t=1}^{\pi} S_t \right) F \right) \end{aligned}$$

Moreover,

$$\sum_{t=1}^{\pi} S_t = Z_{cat} Z_{cat}^T = (U\Sigma V^T)(U\Sigma V^T)^T = U\Sigma^2 U^T.$$

To solve Eq. (6), we let  $Z_{cat}^{F^{(i)}} = [\sqrt{2}Q^{(i)T}Y, \sqrt{\beta}B^{(i)}] \in \mathbb{R}^{n_i \times (k+m)}$ . According to *Proposition 1*, the optimal solution to Eq. (6) is the  $k$  left singular vectors of  $Z_{cat}^{F^{(i)}}$ . This approach naturally reduces the computational cost of updating  $F^{(i)}$  from  $\mathcal{O}(n_i^2 k)$  to  $\mathcal{O}(n_i(k+m)^2)$ .

**Updating  $Y$ :** When  $F^{(i)}$  is fixed, the optimization problem *w.r.t*  $Y$  is

$$\max_{Y^T Y = I} \operatorname{Tr} \left( Y^T \left( \sum_{i=1}^v Q^{(i)} F^{(i)} F^{(i)T} Q^{(i)T} \right) Y \right). \quad (8)$$

Based on *Proposition 1*, we define the formulation  $Z_{cat}^Y = [Q^{(1)} F^{(1)}, \dots, Q^{(i)} F^{(i)}, \dots, Q^{(v)} F^{(v)}] \in \mathbb{R}^{n \times (vk)}$ . Similarly, the optimal solution to Eq. (8) is the  $k$  left singular vectors of  $Z_{cat}^Y$ . Updating  $Y$  requires only  $\mathcal{O}(n(vk)^2)$  time complexity.

The entire optimization procedure for RISE is summarized in Algorithm 1. The code is available at <https://github.com/RISE2025>.

### Algorithm Analysis

**Space complexity.** The storage burden for RISE are the matrices  $B^{(i)} \in \mathbb{R}^{n_i \times m}$ ,  $F^{(i)} \in \mathbb{R}^{n_i \times k}$ , and  $Y \in \mathbb{R}^{n \times k}$ . Thus, the space complexity of RISE is  $\mathcal{O}(nk + n_i k + n_i m)$ .

**Time complexity.** The time complexity of RISE involves updating  $F^{(i)}$  and  $Y$ . Updating  $F^{(i)}$  takes  $\mathcal{O}(n_i(k+m)k)$ ,

Dataset	Size	Views	Classes	Dimensions
Prokaryotic	551	3	4	438,3,393
WebKB	1051	2	2	2949,334
Caltech101-7	1474	6	7	48,40,254,1984,512,928
Wikipedia	2866	2	10	128,10
CIFAR10	50000	3	10	512,2048,1024
FMNIST	60000	3	10	512,512,1280
YoutubeFace	101499	5	31	64,512,64,647,838

Table 1: Descriptions of Multi-view Datasets

while and solving for  $Y$  costs  $\mathcal{O}(nvk^2)$ . Thus, at each iteration, the time complexity is  $\mathcal{O}(\sum_{i=1}^v n_i(k+m)k + nvk^2)$ . After optimization, we perform  $k$ -means with  $\mathcal{O}(nc)$  to output the discrete clustering indicators. Considering that  $k \ll n$ ,  $m \ll n$ , and  $c \ll n$ , the overall time complexity of the proposed optimization algorithm is linear to the number of samples  $n$ . This demonstrates the high efficiency of RISE.

**Convergence.** The optimization algorithm 1 updates  $F^{(i)}$  and  $Y$  by alternatively solving Eq. (6) and Eq. (8). The SVD-based solution designed in *Proposition 1* converges to the global minimum for each sub-problem. Moreover, the objective function in Eq. (5) has a lower bound. These factors collectively ensure the convergence of the algorithm. More details in the appendix. More details in the appendix.

## Experiments

### Experimental Settings

**Datasets and Incomplete Data Construction.** We evaluate our RISE method on seven popular multi-view datasets, which include Prokaryotic (Brbić and Kopriva 2018), Wikipedia (Costa Pereira et al. 2014), WebKB (Bisson and Grimal 2012), Caltech101-7 (Fei-Fei, Fergus, and Perona 2006), CIFAR10 (Krizhevsky, Hinton et al. 2009), FMNIST (Xiao, Rasul, and Vollgraf 2017), and YoutubeFace (Lior Wolf 2011). The details of these datasets are described in Table 1.

Following the approaches in (Wang et al. 2022a; Wen et al. 2023b), we construct the incomplete data with 9 different missing rates  $p = [0.1 : 0.1 : 0.9]$ . For example, when  $p = 0.9$ , we randomly select 10% samples as complete data and randomly drop partial views of the rest 90% samples.

**Compared Methods.** We compare our RISE with ten state-of-the-art IMVC methods. These methods include UEAF (Wen et al. 2019), PIMVC (Deng et al. 2023), HCLS-CGL (Wen et al. 2023a), GSRIMC (Li et al. 2024c), IMVTSC-MVI (Wen et al. 2021b), sFSR-IMVC (Long et al. 2024), DAIMC (Hu and Chen 2019), PSIMVC-PG (Li et al. 2024b), IMVC-CBG (Wang et al. 2022a), and SIMVC-SA (Wen et al. 2023b).

**Implementation Details.** For our RISE method, we use  $K$ -means for anchor selection and the approach in (He et al. 2020) to construct the initial bipartite graph  $\{B^{(i)}\}_{i=1}^v$ . We tune the parameter  $\beta$  within a range of  $[0.01, 0.1, 1, 10, 20, 50, 100, 500, 1000]$ . For the number of anchor points,  $m$  is tuned within a range of  $[1c, 6c]$  for the Prokaryotic, CIFAR10, FMNIST, and YoutubeFace datasets,

Metrics	Methods	UEAF	PIMVC	HCLS-CGL	GSRIMC	IMVTSC-MVI	DAIMC	sFSR-IMVC	PSIMVC-PG	IMVC-CBG	SIMVC-SA	Ours
ACC	Prokaryotic	53.49	54.03	56.62	49.36	53.85	51.31	55.13	47.16	54.91	<u>60.42</u>	<b>75.13</b>
	WebKB	76.56	86.06	81.99	<u>87.30</u>	81.54	81.77	78.72	71.08	86.66	83.60	<b>93.02</b>
	Caltech101-7	37.26	66.88	<u>72.36</u>	65.22	60.53	46.97	67.60	48.52	59.88	59.37	<b>75.25</b>
	Wikipedia	46.01	46.06	39.73	47.52	45.13	46.87	<u>47.98</u>	47.44	47.18	47.22	<b>48.64</b>
	CIFAR10	-	-	-	-	-	92.80	<u>62.77</u>	95.96	95.91	<u>96.04</u>	<b>96.54</b>
	FMNIST	-	-	-	-	-	25.05	10.76	22.20	22.74	<u>50.99</u>	<b>52.49</b>
	YoutubeFace	-	-	-	-	-	-	-	14.32	<u>23.81</u>	15.68	<b>25.52</b>
NMI	Prokaryotic	27.12	25.87	7.28	17.02	19.54	15.98	11.05	22.84	31.25	<u>32.68</u>	<b>38.60</b>
	WebKB	17.24	23.92	17.01	<u>40.91</u>	9.19	16.62	2.25	10.76	34.05	25.14	<b>53.61</b>
	Caltech101-7	27.88	54.09	<b>57.51</b>	53.18	46.52	48.09	44.75	42.80	43.13	40.72	<u>54.06</u>
	Wikipedia	36.11	34.11	28.27	23.75	38.94	32.37	34.36	35.78	<u>37.96</u>	<b>39.18</b>	35.78
	CIFAR10	-	-	-	-	-	89.28	72.79	90.56	90.60	<u>90.67</u>	<b>91.56</b>
	FMNIST	-	-	-	-	-	10.63	0.06	3.11	6.01	<u>27.96</u>	<b>38.73</b>
	YoutubeFace	-	-	-	-	-	-	-	<u>11.54</u>	11.51	8.58	<b>17.88</b>
Purity	Prokaryotic	66.90	63.96	59.58	61.67	58.08	58.41	57.75	59.04	68.81	<u>72.75</u>	<b>77.52</b>
	WebKB	80.90	84.95	82.22	<u>88.50</u>	81.54	81.65	78.74	78.29	86.67	83.60	<b>93.39</b>
	Caltech101-7	77.56	86.94	86.66	58.03	85.48	83.33	74.50	80.61	81.03	79.27	<b>88.59</b>
	Wikipedia	49.39	50.77	44.48	39.10	50.76	49.02	37.17	51.10	50.89	<u>51.58</u>	<b>57.04</b>
	CIFAR10	-	-	-	-	-	93.15	62.81	95.96	95.91	<u>96.04</u>	<b>96.54</b>
	FMNIST	-	-	-	-	-	25.55	10.79	22.20	22.80	<u>51.16</u>	<b>60.17</b>
	YoutubeFace	-	-	-	-	-	-	-	28.11	<u>28.99</u>	28.08	<b>37.78</b>

Table 2: The average ACC, NMI, Purity results of different methods with nine missing ratios on benchmark datasets. “-” indicates out of CPU memory. The best results are highlighted in bold, while the second-best results are marked with underline.

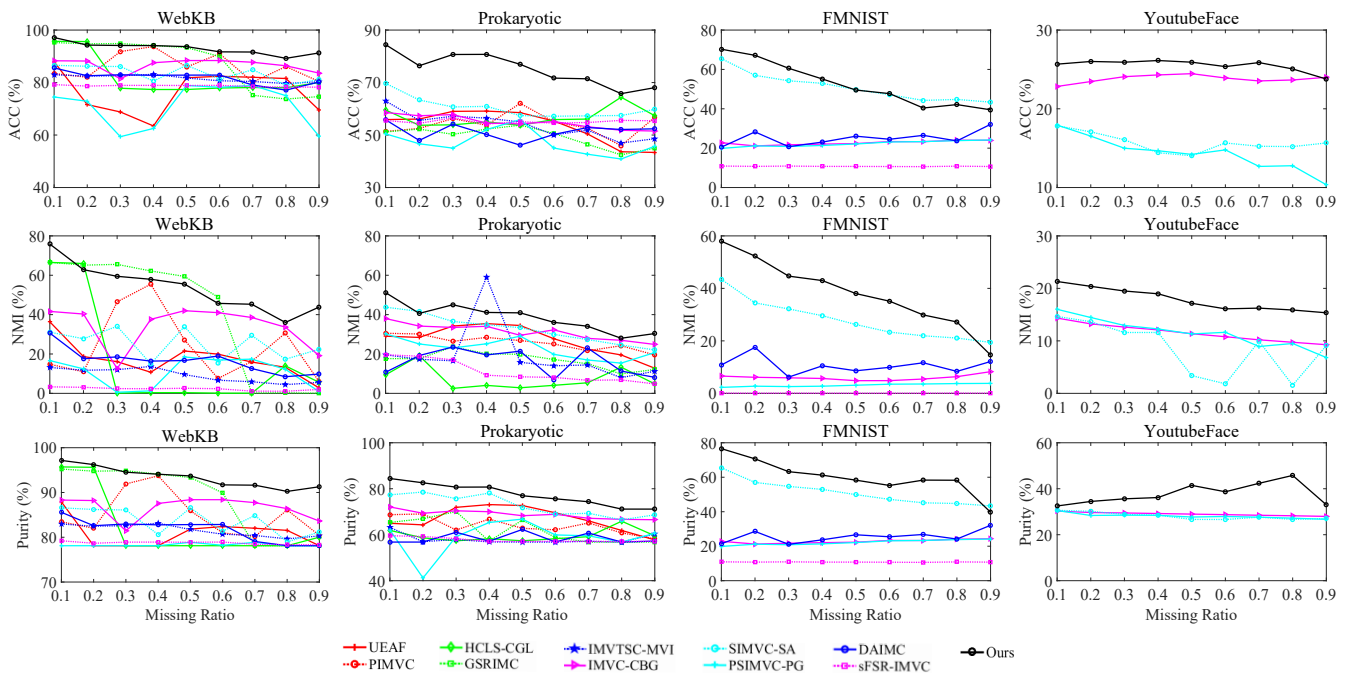


Figure 3: The ACC, NMI and Purity results of different methods with different missing ratios on partial benchmark datasets.

and  $[2c, 10c]$  for the Wikipedia and WebKb datasets. Moreover, the embedded dimension  $k$  is adjusted as an integer multiple of the number of categories and kept smaller than the number of anchors. For all competitors, we obtain their public source codes from open websites and run these algorithms with their default parameter settings. Otherwise, we perform parameter search for better performance. To alleviate the sensitivity of random initialization, we repeat all experiments 10 times and report the average performance. Following previous works (Zhang et al. 2019), clustering accuracy (ACC), normalized mutual information (NMI), and pu-

rity are utilized as metrics to evaluate the performance of all methods. All experiments are performed on a machine with Inter core i7-9700 CPU, 32GB RAM, and Matlab 2022b (64bit).

## Experimental Results and Analysis

Table 2 presents the average clustering results of different methods with nine missing ratios. We can draw the following conclusions.

1. Our proposed RISE achieves the best performance across

Datasets	UEAF	PIMVC	HCLS-CGL	GSRIMC	IMVTSC-MVI	DAIMC	sFSR-IMVC	PSIMVC-PG	IMVC-CBG	SIMVC-SA	Ours
Prokaryotic	99.15	2.94	26.72	128.92	352.19	73.45	7.72	8.71	17.25	16.74	<b>0.37</b>
WebKB	526.35	7.20	21.14	239.32	1024.35	651.86	10.31	13.27	17.08	11.49	<b>0.66</b>
Caltech101-7	635.6	<b>9.82</b>	165.91	623.22	8433.70	1120.39	15.15	28.41	37.87	50.29	<u>11.63</u>
Wikipedia	2344.69	9.37	280.86	2366.09	3916.00	83.10	27.01	28.33	28.92	9.21	<b>2.84</b>
CIFAR10	-	-	-	-	-	93468.10	520.11	596.21	627.45	1359.07	<b>149.99</b>
FMNIST	-	-	-	-	-	89872.00	1602.51	846.43	567.30	973.02	<b>371.89</b>
YoutubeFace	-	-	-	-	-	-	-	1423.55	1866.46	3209.01	<b>899.54</b>

Table 3: Time costs of different methods on benchmark datasets (measured in seconds). “-” means out of CPU memory. The best results are highlighted in bold, while the second-best results are marked with underline.




Figure 4: Clustering results of our RISE with different values of  $\beta$  on two datasets.




Figure 6: The variation of the objective function values for our RISE on two datasets.

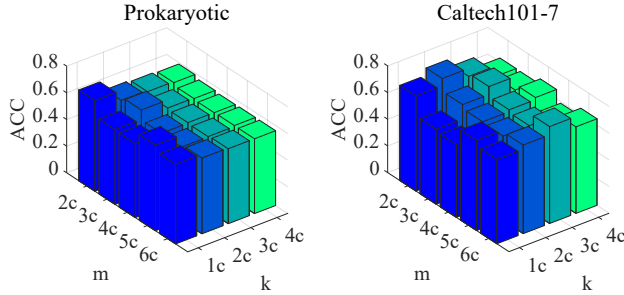


Figure 5: The ACC results of our RISE with different values of  $m$  and  $k$  on two datasets.

the three metrics against all baseline algorithms in most cases, indicating its superiority in IMVC tasks. Notably, SIMVC-SA performs better than other methods. Our RISE surpasses SIMVC-SA with improvements of 14.71%, 9.42%, 15.88%, 1.42%, 0.5%, 1.5%, and 9.84% in terms of ACC on the Prokaryotic, WebKB, Caltech101-7, Wikipedia, CIFAR10, FMNIST, and YoutubeFace datasets, respectively. Similar trends are observed across other metrics, further demonstrating the effectiveness of our RISE, even on large-scale datasets.

2. Compared to instance-level methods (i.e., UEAF, PIMVC, IMVTSC-MVI, sFSR-IMVC, and DAIMC) and graph-level methods (i.e., HCLS-CGL, GSRIMC, PSIMVC-PG, IMVC-CBG, and SIMVC-SA), our RISE based on spectral embedding completion demonstrates its effectiveness.

To provide a more intuitive comparison of different methods across various missing ratios, we present the ACC, NMI, and Purity curves for partial datasets in Fig. 3. As shown in Fig. 3, our RISE outperforms other methods, demonstrating superior and more stable performance. This highlights the

advantage of the second-order rotational-invariant property in multi-view fusion.

Table 3 presents the running times consumed by different methods. As illustrated, our RISE requires the shortest running time on all datasets with the exception of Caltech101-7. This demonstrates the efficiency of RISE. Note that SIMVC-SA and IMVC-CBG take 3209.01 seconds and 1866.46 seconds on the YoutubeFace dataset, while our method only takes 899.54 seconds, reducing the running time by over half. This is because RISE employs bipartite graphs to preserve the essential intrinsic geometric structure of data while reducing redundancy. Moreover, the proposed optimization algorithm further accelerates the clustering process.

## Parameter Sensitivity and Convergence Analysis

To investigate the parameter sensitivity of RISE, we conducted experiments with different parameter settings on two datasets. As shown in Fig. 4, the parameter  $\beta$  indeed affects performance. Specifically, the proposed algorithm remains stable over a large range of  $\beta$ , except when it is less than 10, in which case the NMI metric suffers from a significant degradation. Additionally, we investigated the effect of varying the number of anchors  $m$  and embedded dimensions  $k$  under the optimal setting of other variables. As shown in Fig 5, both the number of anchors and embedded dimension influence the performance of our algorithm. This means that the appropriate values for  $m$  and  $k$  need to be chosen for RISE to obtain the optimal results. Fortunately, only a small amount of tuning is required to achieve satisfactory results.

In addition, we conducted several experiments to validate the convergence of the proposed algorithm. As shown in Fig. 6, the objective function value of our algorithm decreases monotonically in each iteration and converges rapidly. These results verify the convergence of our proposed algorithm.




Figure 7: The ACC results of our RISE with different anchor selection strategies on two datasets.




Figure 8: Clustering results of our RISE with different completion strategies on two datasets.

## Ablation Study

**Effect of Anchor Selection Strategy.** To investigate the effect of different strategies for generating anchors, we used the k-means, random selection, and directly alternate sampling (DAS) method (Li et al. 2020) for anchor selection. As shown in Fig. 7, the k-means is more suitable for our algorithm. Specifically, the  $k$ -means based anchors achieve over 15% performance improvement compared to the DAS approach on the Prokaryotic dataset, and around 10% improvement on the Caltech101-7 dataset. This is likely because the DAS approach is designed for complete multi-view data, while  $k$ -means is a more general method.

**Performance Effect of Completion Strategy.** The rotation-invariant complete representation learning strategy is the main contribution of this paper. To further demonstrate its effectiveness, we conducted ablation experiments comparing different complete representation learning strategies. The term 'metric-2' refers to our proposed method, while the term 'metric-1' indicates the use of Eq. (3) to learn the complete representation. As shown in Fig. 8, the proposed rotation-invariant learning module effectively improves the algorithm's performance on the Prokaryotic and Caltech101-7 datasets in terms of ACC and NMI. This demonstrates the effectiveness of our proposed strategy. Additionally, our strategy shows smaller performance fluctuations as the missing rate increases, indicating the superiority of our approach.

**Efficiency Effect of Completion Strategy.** The proposed optimization algorithm enhances the efficiency of our RISE method. To further demonstrate its effectiveness, we plot the execution time comparison histograms of learning strategies using Eq. (3) and our proposed second-order rotation-invariant complete representation learning in Fig. 9. From this, we observe that our RISE method requires less time




Figure 9: Time costs of our RISE with different completion strategies on two datasets.




Figure 10: Clustering results with bipartite graph and full-size similarity graph on two datasets.

compared to the first-order learning strategy using Eq. (3). Specifically, our RISE shows nearly 10 times lower execution time on the Prokaryotic dataset and around 100 times lower on the Caltech101-7 dataset.

**Effect of Graph Selection.** To evaluate the performance of bipartite graph in reducing data redundancy, we used both bipartite graph and full similarity graph to separately describe the neighborhood structure among data. We then executed the rotation-invariant spectral embedding framework to obtain the consensus complete representation and applied  $K$ -means to get the final clustering indicators. As is shown in Fig. 10, the bipartite graph can provide more superior performance compared to the full similarity graph. Additionally, using full similarity graph results in higher optimization costs because the proposed simplified SVD-based optimization strategy cannot be used.

## Conclusions

This paper proposed a highly efficient incomplete multi-view learning framework, called RISE. It developed a fast late spectral embedding completion technique. RISE learned view-specific embeddings from incomplete bipartite graphs and then recovered the complete consensus representation from these embeddings using the proposed second-order rotation-invariant learning module. Both processes were integrated into a unified framework to mutually enhance performance. To solve the proposed objective function, we designed a fast two-step alternating optimization algorithm based on double SVD decompositions. Extensive experiments on datasets of varying sizes demonstrated the effectiveness, scalability, and efficiency of our proposed method. In the future, developing a parameter-free model will require more attention.

## Acknowledgments

This work was funded by the Science and Technology Development Fund, Macau SAR (File no. 0049/2022/A1, 0050/2024/AGJ), by the University of Macau Development Foundation (File no. MYRG-GRG2024-00181-FST-UMDF).

## References

- Bisson, G.; and Grimal, C. 2012. Apprentissage multi-vue de co-similarit s pour la classification. In *Conf rence Francophone sur l'Apprentissage Automatique*.
- Brbić, M.; and Kopriva, I. 2018. Multi-view low-rank sparse subspace clustering. *Pattern recognition*, 73: 247–258.
- Chen, M.-S.; Huang, L.; Wang, C.-D.; and Huang, D. 2020. Multi-view clustering in latent embedding space. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 3513–3520.
- Chen, Y.; Zhao, X.; Zhang, Z.; Liu, Y.; Su, J.; and Zhou, Y. 2023. Tensor learning meets dynamic anchor learning: From complete to incomplete multiview clustering. *IEEE Transactions on Neural Networks and Learning Systems*.
- Costa Pereira, J.; Coviello, E.; Doyle, G.; Rasiwasia, N.; Lanckriet, G.; Levy, R.; and Vasconcelos, N. 2014. On the role of Correlation and Abstraction in Cross-Modal Multimedia Retrieval. *Transactions of Pattern Analysis and Machine Intelligence*, 36(3): 521–535.
- Cui, J.; Fu, Y.; Huang, C.; and Wen, J. 2022. Low-rank graph completion-based incomplete multiview clustering. *IEEE Transactions on Neural Networks and Learning Systems*.
- Deng, S.; Wen, J.; Liu, C.; Yan, K.; Xu, G.; and Xu, Y. 2023. Projective incomplete multi-view clustering. *IEEE Transactions on Neural Networks and Learning Systems*.
- Fei-Fei, L.; Fergus, R.; and Perona, P. 2006. One-shot learning of object categories. *IEEE transactions on pattern analysis and machine intelligence*, 28(4): 594–611.
- Gao, H.; Peng, Y.; and Jian, S. 2016. Incomplete multi-view clustering. In *Intelligent Information Processing VIII: 9th IFIP TC 12 International Conference, IIP 2016, Melbourne, VIC, Australia, November 18-21, 2016, Proceedings 9*, 245–255. Springer.
- Hao, W.; Pang, S.; Bai, X.; and Xue, J. 2023. Tensor-Based Incomplete Multi-View Clustering With Low-Rank Data Reconstruction and Consistency Guidance. *IEEE Transactions on Circuits and Systems for Video Technology*.
- He, F.; Nie, F.; Wang, R.; Hu, H.; Jia, W.; and Li, X. 2020. Fast semi-supervised learning with optimal bipartite graph. *IEEE Transactions on Knowledge and Data Engineering*, 33(9): 3245–3257.
- Hu, M.; and Chen, S. 2019. Doubly aligned incomplete multi-view clustering. *arXiv preprint arXiv:1903.02785*.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Li, J.; Gao, Q.; Wang, Q.; and Xia, W. 2024a. Tensorized Label Learning on Anchor Graph. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 13537–13544.
- Li, M.; Wang, S.; Liu, X.; and Liu, S. 2024b. Parameter-free and scalable incomplete multiview clustering with prototype graph. *IEEE Transactions on Neural Networks and Learning Systems*, 35(1): 300–310.
- Li, X.; Sun, Y.; Sun, Q.; Dai, J.; and Ren, Z. 2023. Distribution Consistency based Fast Anchor Imputation for Incomplete Multi-view Clustering. In *Proceedings of the 31st ACM International Conference on Multimedia*, 368–376.
- 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.
- Li, X.-L.; Chen, M.-S.; Wang, C.-D.; and Lai, J.-H. 2024c. Refining graph structure for incomplete multi-view clustering. *IEEE transactions on neural networks and learning systems*.
- Li, Z.; Tang, C.; Zheng, X.; Liu, X.; Zhang, W.; and Zhu, E. 2022. High-order correlation preserved incomplete multi-view subspace clustering. *IEEE Transactions on Image Processing*, 31: 2067–2080.
- Lin, Y.; Gou, Y.; Liu, Z.; Li, B.; Lv, J.; and Peng, X. 2021. Completer: Incomplete multi-view clustering via contrastive prediction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11174–11183.
- Lior Wolf, I. M., Tal Hassner. 2011. "YouTube Faces DB" Technical Report.
- Liu, C.; Wu, S.; Li, R.; Jiang, D.; and Wong, H.-S. 2023. Self-supervised graph completion for incomplete multi-view clustering. *IEEE Transactions on Knowledge and Data Engineering*.
- Liu, J.; Liu, X.; Zhang, Y.; Zhang, P.; Tu, W.; Wang, S.; Zhou, S.; Liang, W.; Wang, S.; and Yang, Y. 2021. Self-representation subspace clustering for incomplete multi-view data. In *Proceedings of the 29th ACM international conference on multimedia*, 2726–2734.
- Liu, S.; Liu, X.; Wang, S.; Niu, X.; and Zhu, E. 2022a. Fast incomplete multi-view clustering with view-independent anchors. *IEEE Transactions on Neural Networks and Learning Systems*.
- Liu, S.; Wang, S.; Zhang, P.; Xu, K.; Liu, X.; Zhang, C.; and Gao, F. 2022b. Efficient one-pass multi-view subspace clustering with consensus anchors. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 7576–7584.
- Long, Z.; Zhu, C.; Comon, P.; Ren, Y.; and Liu, Y. 2024. Feature Space Recovery for Efficient Incomplete Multi-View Clustering. *IEEE Transactions on Knowledge and Data Engineering*, 36(9): 4664–4677.
- Lv, Z.; Gao, Q.; Zhang, X.; Li, Q.; and Yang, M. 2022. View-consistency learning for incomplete multiview clustering. *IEEE Transactions on Image Processing*, 31: 4790–4802.
- Shao, W.; He, L.; and Yu, P. S. 2015. Multiple incomplete views clustering via weighted nonnegative matrix factorization with regularization. In *Joint European conference on machine learning and knowledge discovery in databases*, 318–334. Springer.

- Shen, Q.; Xu, T.; Liang, Y.; Chen, Y.; and He, Z. 2023. Robust Tensor Recovery for Incomplete Multi-view Clustering. *IEEE Transactions on Multimedia*.
- Von Luxburg, U. 2007. A tutorial on spectral clustering. *Statistics and computing*, 17: 395–416.
- Wang, S.; Liu, X.; Liu, L.; Tu, W.; Zhu, X.; Liu, J.; Zhou, S.; and Zhu, E. 2022a. Highly-efficient incomplete large-scale multi-view clustering with consensus bipartite graph. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9776–9785.
- Wang, S.; Liu, X.; Liu, S.; Jin, J.; Tu, W.; Zhu, X.; and Zhu, E. 2022b. Align then fusion: Generalized large-scale multi-view clustering with anchor matching correspondences. *Advances in Neural Information Processing Systems*, 35: 5882–5895.
- Wen, J.; Liu, C.; Xu, G.; Wu, Z.; Huang, C.; Fei, L.; and Xu, Y. 2023a. Highly confident local structure based consensus graph learning for incomplete multi-view clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15712–15721.
- Wen, J.; Xu, Y.; and Liu, H. 2020. Incomplete multiview spectral clustering with adaptive graph learning. *IEEE transactions on cybernetics*, 50(4): 1418–1429.
- Wen, J.; Yan, K.; Zhang, Z.; Xu, Y.; Wang, J.; Fei, L.; and Zhang, B. 2020. Adaptive graph completion based incomplete multi-view clustering. *IEEE Transactions on Multimedia*, 23: 2493–2504.
- Wen, J.; Zhang, Z.; Xu, Y.; Zhang, B.; Fei, L.; and Liu, H. 2019. Unified embedding alignment with missing views inferring for incomplete multi-view clustering. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 5393–5400.
- Wen, J.; Zhang, Z.; Zhang, Z.; Fei, L.; and Wang, M. 2021a. Generalized incomplete multiview clustering with flexible locality structure diffusion. *IEEE transactions on cybernetics*, 51(1): 101–114.
- Wen, J.; Zhang, Z.; Zhang, Z.; Zhu, L.; Fei, L.; Zhang, B.; and Xu, Y. 2021b. Unified tensor framework for incomplete multi-view clustering and missing-view inferring. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 10273–10281.
- Wen, Y.; Wang, S.; Liang, K.; Liang, W.; Wan, X.; Liu, X.; Liu, S.; Liu, J.; and Zhu, E. 2023b. Scalable Incomplete Multi-View Clustering with Structure Alignment. In *Proceedings of the 31st ACM International Conference on Multimedia*, 3031–3040.
- Xia, D.; Yang, Y.; and Yang, S. 2023. Incomplete Multi-View Clustering via Auto-Weighted Fusion in Partition Space. *Tsinghua Science and Technology*, 28(3): 595–611.
- Xia, W.; Gao, Q.; Wang, Q.; and Gao, X. 2022. Tensor completion-based incomplete multiview clustering. *IEEE Transactions on Cybernetics*, 52(12): 13635–13644.
- Xiao, H.; Rasul, K.; and Vollgraf, R. 2017. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*.
- Yin, J.; and Sun, S. 2023. Incomplete multi-view clustering with reconstructed views. *IEEE Transactions on Knowledge and Data Engineering*, 35(3): 2671–2682.
- Zhang, C.; Fu, H.; Hu, Q.; Cao, X.; Xie, Y.; Tao, D.; and Xu, D. 2018. Generalized latent multi-view subspace clustering. *IEEE transactions on pattern analysis and machine intelligence*, 42(1): 86–99.
- Zhang, C.; Li, H.; Lv, W.; Huang, Z.; Gao, Y.; and Chen, C. 2023a. Enhanced tensor low-rank and sparse representation recovery for incomplete multi-view clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 11174–11182.
- Zhang, C.; Wei, J.; Wang, B.; Li, Z.; Chen, C.; and Li, H. 2023b. Robust Spectral Embedding Completion Based Incomplete Multi-view Clustering. In *Proceedings of the 31st ACM International Conference on Multimedia*, 300–308.
- Zhang, G.-Y.; Huang, D.; and Wang, C.-D. 2024. Unified and Tensorized Incomplete Multi-View Kernel Subspace Clustering. *IEEE Transactions on Emerging Topics in Computational Intelligence*.
- Zhang, Y.; Wu, J.; Zhou, C.; Cai, Z.; Yang, J.; and Yu, P. S. 2019. Multi-view fusion with extreme learning machine for clustering. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(5): 1–23.
- Zhao, S.; Wen, J.; Fei, L.; and Zhang, B. 2023a. Tensorized incomplete multi-view clustering with intrinsic graph completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 11327–11335.
- Zhao, X.; Shen, Q.; Chen, Y.; Liang, Y.; Chen, J.; and Zhou, Y. 2023b. Self-completed bipartite graph learning for fast incomplete multi-view clustering. *IEEE Transactions on Circuits and Systems for Video Technology*.